

Efficient Empirical Revenue Maximization in Single-Parameter Auction Environments

Yannai A. Gonczarowski*

Noam Nisan†

October 31, 2016

Abstract

We present a polynomial-time algorithm that, given samples from the unknown valuation distribution of each bidder, learns an auction that approximately maximizes the auctioneer’s revenue in a variety of single-parameter auction environments including matroid environments, position environments, and the public project environment. The valuation distributions may be arbitrary bounded distributions (in particular, they may be irregular, and may differ for the various bidders), thus resolving a problem left open by previous papers. The analysis uses basic tools, is performed in its entirety in value-space, and simplifies the analysis of previously known results for special cases. Furthermore, the analysis extends to certain single-parameter auction environments where precise revenue maximization is known to be intractable, such as knapsack environments.

1 Introduction

We start by considering auctions of a single item to n bidders (we later extend our results beyond the single-item auction environment). The value (i.e., maximum willingness to pay) v_i for the item of each bidder i is distributed according to some *unknown* (perhaps irregular) distribution F_i that is supported on a known bounded interval $[0, H]$, and the values of the bidders are drawn independently of each other. An auctioneer is given $t = \text{poly}(H, n, 1/\varepsilon)$ random samples from each F_i , and aims to design an auction that approximately maximizes her revenue from (future bidders with values drawn from) $F_1 \times \cdots \times F_n$.

The most natural approach would be to take, for each bidder i , the empirical distribution \hat{F}_i defined as the uniform distribution over the (multi)set of the t samples from F_i , and design an auction that maximizes the empirical revenue, i.e., the revenue from $\hat{F}_1 \times \cdots \times \hat{F}_n$. One would hope that the resulting auction would approximately maximize the revenue from $F_1 \times \cdots \times F_n$ as well.

*Einstein Institute of Mathematics, Rachel & Selim Benin School of Computer Science & Engineering and Federmann Center for the Study of Rationality, The Hebrew University of Jerusalem, Israel; and Microsoft Research, *E-mail*: yannai@gonch.name.

†Rachel & Selim Benin School of Computer Science & Engineering and Federmann Center for the Study of Rationality, The Hebrew University of Jerusalem, Israel; and Microsoft Research, *E-mail*: noam.nisan@gmail.com.

Unfortunately, such an approach is prone to overfitting the samples, in which case the empirical-revenue-maximizing auction may perform poorly on $F_1 \times \cdots \times F_n$. Therefore, recent papers have focused on designing auctions that approximately maximize the empirical revenue, in a way that is robust against overfitting. These papers combine the delicate understanding of the nature of optimal auctions due to Myerson (1981) with various techniques of converting an auction that is good for variants of the empirical distribution to one that is good for the original distribution, in several special cases: a single buyer with a regular value distribution (Huang et al., 2015), multiple buyers with regular distributions (Cole and Roughgarden, 2014; Devanur et al., 2016), and multiple buyers with i.i.d. irregular distributions (Roughgarden and Schrijvers, 2016). The general case (i.e., that of non-i.i.d., possibly irregular, distributions) was left open, even though it was shown that, information-theoretically, it should be possible (Morgenstern and Roughgarden, 2015).

In this paper, we resolve the above-described open problem by showing how to take the auction that is optimal for the empirical sampled distribution $\hat{F}_1 \times \cdots \times \hat{F}_n$, and efficiently “round” it to integer multiples of ε in a way that ensures that the rounded auction is near-optimal also for the original distribution $F_1 \times \cdots \times F_n$. The analysis uses basic tools, is performed in its entirety in value-space (rather than quantile-space, or “virtual-value”-space), and also simplifies the analysis of above-described previously known results for special cases. We now give a high-level review of our analysis for the basic case of a single-item auction. In Section 6, we extend our analysis to a variety of single-parameter auction environments including matroid environments, position environments, and the public project environment, as well as to certain single-parameter auction environments where precise revenue maximization is known to be intractable, such as knapsack environments.

We start by defining what we mean by “rounding” an auction. To do so, we view the domain $[0, H]$ of the possible valuations (values) of each bidder as composed of the disjoint union of the semiopen ε -intervals $[j \cdot \varepsilon, (j + 1) \cdot \varepsilon)$ for $j = 0, \dots, H/\varepsilon$.

Definition 1.1 (ε -Coarse Auction; ε -Rounding).

- We call an auction ε -coarse if its *outcome*, i.e., (item) *allocation* and payment, is constant within each product of ε -intervals. That is, the outcome of an ε -coarse auction for bids v_1, \dots, v_n depends only on the indices j_1, \dots, j_n of the ε -intervals to which the bids v_1, \dots, v_n respectively belong (i.e., $j_i \cdot \varepsilon \leq v_i < (j_i + 1) \cdot \varepsilon$ for every i).
- We ε -round an auction A to an ε -coarse auction A' by choosing for each ε -interval j of each bidder i a fixed value $v_j^i \in [j \cdot \varepsilon, (j + 1) \cdot \varepsilon)$ in that ε -interval, and defining the allocation rule of the auction A' to treat all bids of bidder i in that ε -interval as the auction A treats v_j^i . In other words, A' is the unique ε -coarse auction whose allocation on every tuple $(v_{j_1}^1, \dots, v_{j_n}^n)$ of these chosen values is identical to that of A . The payments of A' are defined to be the unique payments that make it incentive compatible, and are thus generally *not* identical to those of A .

Our notion of rounding auctions allows for many different ways of ε -rounding any given auction, depending on the choice of the values v_j^i . One might hope that all of these roundings achieve approximately the same revenue as the original auction, but that is

not the case.¹ Our main lemma shows that *some* ε -rounding of any given auction does achieve approximately the revenue of the unrounded auction. Moreover, for an auction that is optimal for the product of uniform distributions over finite sets of (sampled) values, finding such an ε -rounding can be done in polynomial time.

Lemma 1.2 (See also Propositions 3.7 and 4.3).

- a. For every product distribution $F_1 \times \dots \times F_n$ and auction A , there exists an ε -coarse auction A' that is an ε -rounding of A and whose revenue from $F_1 \times \dots \times F_n$ is less than an additive ε smaller than that of A .*
- b. There exists a deterministic algorithm that runs in time $\text{poly}(H, n, 1/\varepsilon, t)$, takes as input i) a product distribution $\hat{F}_1 \times \dots \times \hat{F}_n$ where each \hat{F}_i is uniform over a multiset of t (sampled) values, and ii) an auction A that is optimal for $\hat{F}_1 \times \dots \times \hat{F}_n$, and outputs an ε -coarse auction A' that is an ε -rounding of A and whose revenue from $\hat{F}_1 \times \dots \times \hat{F}_n$ is less than an additive $n\varepsilon$ smaller than that of A .*

As opposed to existing auction-rounding schemes in the literature (Hart and Nisan, 2013; Dughmi et al., 2014), the rounding in Lemma 1.2 is not “universal” for all distributions (i.e., the rounded auction A' does not approximate the revenue of A on all product distributions), but rather needs to be tailored specifically for $F_1 \times \dots \times F_n$.² Furthermore, as opposed to Devanur et al. (2016), we round the optimal auction for the empirical distribution, rather than finding the optimal auction for a discretized version of the empirical distribution; this allows us to handle unrestricted bounded distributions, without any assumption such as regularity. Our algorithm considers a randomized rounding, where all values $j \cdot \varepsilon \leq v_i < (j+1) \cdot \varepsilon$ are rounded to a fixed value v_j^i that is itself randomly chosen according to the distribution F_i conditioned on being in this ε -interval. We show that the *expected* revenue of the resulting auction (or more accurately, distribution over auctions) is within less than an additive ε of the revenue of A , and so some deterministic auction within this distribution over auctions loses less than an additive ε in revenue compared to A . When for every $i \in N$ it is the case that $F_i = \hat{F}_i$ is a uniform distribution over a multiset of t (sampled) values, our algorithm deterministically and efficiently searches for a deterministic auction that loses less than an additive $n\varepsilon$, within the above-defined randomized distribution over possible roundings.

Once we have Lemma 1.2, the algorithm for empirical revenue maximization is simple: Let A be Myerson’s optimal auction for the empirically sampled product distribution $\hat{F}_1 \times \dots \times \hat{F}_n$ (which is readily computed from the samples as spelled out in Elkind, 2007), and output the ε -rounded approximately optimal auction A' produced by the rounding

¹While it is true that the price paid by a bidder i for winning the item in a rounded auction is indeed approximately the same whether or not we round all bids of that bidder, there are two complications: the minor one is that bids of bidder i that used to win the item in the original auction may cease to do so in the rounded auction; the major one is that in the original auction, tiny changes in the bid of bidder i may result in arbitrary changes in the payment of another bidder j .

²We show that this is unavoidable by giving an example of a two-bidder auction that is optimal for some product of regular distributions and yet for every $\varepsilon > 0$, every ε -coarse auction must lose $\Omega(1)$ (i.e., at least some constant independent of ε) revenue for some valuation profile (v_1, v_2) . This is in contrast to the case of i.i.d. distributions, for which we show in Appendix C how to ε -round any auction that is optimal for some $F^n = F \times \dots \times F$ into a single ε -coarse auction that loses less than an additive ε in revenue universally for every profile of valuations (v_1, \dots, v_n) , and thus also for every distribution. This implies a simplified proof for the recent result of Roughgarden and Schrijvers (2016) for the i.i.d. case, and highlights the difference between the general case considered here and the i.i.d. case.

algorithm for A , tailored to lose less than $n\varepsilon$ revenue on $\hat{F}_1 \times \cdots \times \hat{F}_n$. When the ε -rounding algorithm is applied to Myerson’s (ironed-virtual-welfare maximizing) optimal auction A , the resulting ε -coarse auction A' turns out to be of the following simple form:

Definition 1.3 ((H, ε) -simple auction). We call an auction (H, ε) -simple if there exists a sequence P of distinct pairs of the form (i, j) where $i \in \{1, \dots, n\}$ and $j \in \{0, \dots, \lfloor H/\varepsilon \rfloor\}$, such that the winner is always the first bidder i in the order P that has $v_i \geq j \cdot \varepsilon$, and the winner’s payment is always the threshold value required to win.³

As (H, ε) -simple auctions have a concise polynomial-length description (the sequence P), basic learning-theory intuition implies that overfitting can be ruled out. Specifically, we apply an elegant concentration inequality due to Devanur et al. (2016), which shows that the revenue of any fixed auction from the product of the empirically sampled distributions well approximates its revenue from the product of the true distributions. We then conclude by taking a union bound over the small number of (H, ε) -simple auctions to obtain our main theorem:

Theorem 1.4 (See also Theorem 5.6). *There exists a deterministic algorithm that runs in time $\text{poly}(H, n, 1/\varepsilon, \log 1/\delta)$, takes as input $t = \text{poly}(H, n, 1/\varepsilon, \log 1/\delta)$ random samples from unknown distributions F_1, \dots, F_n supported on $[0, H]$ and, with probability at least $1 - \delta$, outputs a description of an n -bidder single-item auction that maximizes the revenue from $F_1 \times \cdots \times F_n$ up to less than an additive ε . The produced auction is an $(\frac{\varepsilon}{n+2}, H)$ -simple $\frac{\varepsilon}{n+2}$ -rounding of the auction that maximizes revenue from the empirical product of the uniform distributions over the samples.*

As Myerson’s characterization of optimal auctions applies to general *single-parameter* auction environments, we are able to generalize Theorem 1.4 to a wide variety of computationally tractable single-parameter environments by employing similar, yet somewhat more delicate, analysis. First, our main ingredient, Lemma 1.2(a), readily generalizes to arbitrary single-parameter environments, with the error term multiplied by the expected “number of winners”. Nevertheless, it is not possible to generalize Lemma 1.2(b) since the underlying algorithm requires the ability to compute the expected overall revenue of an auction, which may be computationally intractable even for computationally tractable environments. Therefore, we instead take a different approach to prove an analogue of Lemma 1.2(b), by applying random sampling and then derandomizing the process using random bits obtained from the order in which (polynomially many) samples are drawn. Finally, while bounding the number of resulting ε -coarse auctions is easy in some environments, a more sophisticated argument based on Cramer’s rule is required more generally.⁴ Putting all of these together, we derive the following generalization of Theorem 1.4:

Theorem 1.5 (See also Theorem 6.10 and Remark 6.11). *For a class of single-parameter auction environments including computationally tractable deterministic environments (such as matroid environments and the public project environment) and position environments, there exists a deterministic algorithm that runs in time $\text{poly}(H, n, 1/\varepsilon, \log 1/\delta)$, takes*

³ This may be viewed as an ε -discretized variant of the “leveled auctions” of Morgenstern and Roughgarden (2015), albeit with more freedom in tie breaking; this finite discretization is exactly what allows us to provide an efficient algorithm rather than an information-theoretic result. The way we bound the number of (H, ε) -simple auctions is reminiscent of the way in which Devanur et al. (2016) bound the number of optimal auctions on a finite valuation space.

⁴We are not aware of the use of any similar argument in the literature on mechanism design, and hope that this type of argument may find additional uses in similar contexts in the future.

as input $t = \text{poly}(H, n, 1/\varepsilon, \log 1/\delta)$ random samples from unknown distributions F_1, \dots, F_n supported on $[0, H]$ and, with probability at least $1-\delta$, outputs a description of an n -bidder auction for the given auction environment that maximizes the revenue from $F_1 \times \dots \times F_n$ up to less than an additive ε .

Finally, for single-parameter auction environments where precise revenue maximization is known to be intractable, but where efficient approximate revenue maximization up to some multiplicative factor C is possible, we prove, under certain assumptions on this up-to- C maximization algorithm (that are satisfied, e.g., in the case of knapsack environments), a generalization of Theorem 6.10, providing a deterministic polynomial-time algorithm that learns from polynomially many samples a *tractable* auction that with high probability approximates the maximal revenue from $F_1 \times \dots \times F_n$ up to the same multiplicative factor of C , plus less than an additive ε . (See also Theorem 6.14.)

1.1 Structure

The remainder of this paper is structured as follows. Section 2 provides definitions and some background, Section 3 provides the derivation of Lemma 1.2(a), Section 4 provides the derivation of Lemma 1.2(b), and Section 5 provides the remainder of the derivation of Theorem 1.4. In Section 6, we extend the above single-item analysis to more general single-parameter auction environments, and derive Theorem 1.5 and its generalization for single-parameter auction environments where precise revenue maximization is intractable. Some examples referenced in this paper are relegated to Appendix A, and the proofs of some of the results stated throughout this paper are relegated to Appendix B. As mentioned above, Appendix C presents an analogous (yet simpler) analysis for the special case of i.i.d. distributions, where distribution-independent rounding of optimal auctions is possible.

2 Model, Definitions, and Background

2.1 Auctions

The auctions we consider are deterministic direct-revelation dominant-strategy incentive-compatible (DSIC) and ex-post individually rational (IR) single-item auctions among n bidders numbered $N \triangleq \{1, 2, \dots, n\}$, each having a valuation (value) in $[0, H]$ for some $H \in \mathbb{R}_+$ that is known to the auctioneer. We denote by $r^A(v_1, \dots, v_n)$ the revenue of auction A when the bidders have values v_1, \dots, v_n , and for every $i \in N$, denote by $r_i^A(v_1, \dots, v_n)$ the revenue of auction A from bidder i (i.e., the payment of bidder i) when the bidders have values v_1, \dots, v_n . (So $r^A(v_1, \dots, v_n) = \sum_i r_i^A(v_1, \dots, v_n)$.) We furthermore denote by $r_i^A(v_i) \triangleq \mathbb{E}_{\forall j \neq i: v_j \sim F_j} r_i^A(v_1, \dots, v_n)$ the *expected* revenue of auction A from bidder i , when bidder i has value v_i and the valuations of all other bidders j are independently drawn from the respective distributions F_j , which will be clear from context. The overall revenue of auction A from a product distribution $F_1 \times \dots \times F_n$ is denoted by $\text{Rev}^A(F_1 \times \dots \times F_n) \triangleq \mathbb{E}_{v_1 \sim F_1, \dots, v_n \sim F_n} r^A(v_1, \dots, v_n)$.

2.2 Rounding

We now make Definition 1.1 from Section 1 precise, and provide some supporting notation.

Definition 2.1 (ε -Interval). Let $\varepsilon > 0$. An ε -interval is a semiopen interval of the form $[j \cdot \varepsilon, (j + 1) \cdot \varepsilon)$ for some integer $j \in \{0, \dots, \lfloor H/\varepsilon \rfloor\}$. We say that j is the *index* of this ε -interval.

Definition 2.2 (ε -Coarse Auction). Let $\varepsilon > 0$. An n -bidder auction A (for valuations in $[0, H]$) is said to be ε -coarse, if for every pair of valuations profiles $(v_1, \dots, v_n), (w_1, \dots, w_n) \in [0, H]^n$ such that for every $i \in N$ there exists j_i such that $v_i, w_i \in [j_i \cdot \varepsilon, (j_i + 1) \cdot \varepsilon)$, the outcome (allocation and payment) of A is the same for (v_1, \dots, v_n) and for (w_1, \dots, w_n) .

We next define how auctions can be “rounded” into coarse ones, using a sequence of “rounding actions.”

Definition 2.3 (ε -Rounding). Let A be an n -bidder auction and let $\varepsilon > 0$.

- An ε -rounding action is a triplet (i, j, v_j^i) where $i \in N$ is a bidder, j is an index of an ε -interval, and $v_j^i \in [j \cdot \varepsilon, (j + 1) \cdot \varepsilon)$ is a value in that ε -interval. Applying the ε -rounding action (i, j, v_j^i) to the auction A yields an auction A' having the following allocation rule: given bids v_1, \dots, v_n , if $v_i \in [j \cdot \varepsilon, (j + 1) \cdot \varepsilon)$, then change the bid of v_i to v_j^i (for other bidders, or for the i th bidder if her value is not in this ε -interval, keep the original bid unchanged), and then run the allocation rule of A on the (possibly) updated values. The payment rule of A' is *not* directly taken from A , but is rather what is needed to ensure truthfulness: each winning bidder pays her *minimal winning bid* (i.e., the infimum of all bids that would have still allowed her to win when the bids of all other bidders are unchanged).
- An ε -rounding rule is a collection of ε -rounding actions (i, j, v_j^i) , one action for each pair of bidder $i \in N$ and index $j \in \{0, \dots, \lfloor H/\varepsilon \rfloor\}$ of an ε -interval. Applying an ε -rounding rule to the auction A yields an auction A' obtained by applying all ε -rounding actions in the rule, in arbitrary order, to A . I.e., given bids v_1, \dots, v_n , for every bidder i the bid v_i is changed to v_j^i where j is the index of the ε -interval that contains the original bid v_i . Each winning bidder pays her minimal winning bid. Such an auction A' is called an ε -rounding of A .

Remark 2.4. Any ε -rounding of any auction is ε -coarse.

2.3 Optimal Auctions

Our analysis makes very weak use of Myerson’s (1981) characterization of optimal single-item auctions, which we therefore only present to the extent required by our analysis. (We emphasize this weak use of Myerson’s characterization of optimal single-item auctions by using the nonstandard term “Myersonian Auction.”)

Definition 2.5 (Myersonian Auction, Myerson, 1981). An n -bidder *Myersonian auction* (for valuations in $[0, H]$) is a tuple $(\phi_i)_{i \in N}$, where for every $i \in N$, $\phi_i : [0, H] \rightarrow \mathbb{R}$ is a nondecreasing function called the *ironed virtual valuation* of bidder i . In this auction, there is a winner unless $\phi_i(v_i) < 0$ for all $i \in N$, and the winner is the bidder with lowest index among those whose bid v_i maximizes $\phi_i(v_i)$; the winner pays her minimal winning bid.⁵

⁵This auction can also be made symmetric, by choosing the winner uniformly at random between all bidders whose bid v_i maximizes $\phi_i(v_i)$, and adapting the payments accordingly. The revenue is the same either way. In this paper we use lexicographic ordering for simplicity.

Myerson (1981) proved that for every continuous product distribution, there exists a Myersonian auction that obtains the optimal revenue. Elkind (2007) showed the same for discrete product distributions, giving an efficient algorithm for computing this optimal auction.

Theorem 2.6 (Myerson, 1981). *For every product distribution $F = F_1 \times \cdots \times F_n$, there exists a Myersonian auction $(\phi_i)_{i \in N}$, denoted $OPT(F)$, that achieves maximal revenue from F among all possible auctions. Moreover, for every $i \in N$, the ironed virtual valuation ϕ_i depends only on F_i .*

Theorem 2.7 (Elkind, 2007). *Let $t \in \mathbb{N}$. There exists an algorithm that runs in time $\text{poly}(t)$, such that given a discrete distribution \hat{F}_i with support of size at most t , outputs a nondecreasing function $\phi_i : \text{supp } \hat{F}_i \rightarrow \mathbb{R}$ (so, ϕ_i is a nondecreasing sequence of at most t real numbers), such that for every product $\hat{F} = \hat{F}_1 \times \cdots \times \hat{F}_n$ of discrete distributions each having support of size at most t , the Myersonian auction $(\phi_i)_{i \in N}$ (where ϕ_i is the output of the algorithm given \hat{F}_i) achieves maximal revenue from \hat{F} among all possible auctions.*

Remark 2.8. The algorithm of Theorem 2.7 outputs an ironed virtual valuation ϕ_i whose domain is $\text{supp } \hat{F}_i$ rather than $[0, H]$. To be completely formal and avoid any ambiguities with regard to the definition of, e.g., the minimal winning bid of any bidder (and later on, with regard to the revenue of this auction from distributions with support larger than that of \hat{F}), we emphasize that whenever we consider $(\phi_i)_{i \in N}$ as a Myersonian auction, we do so by interpreting each ϕ_i as specifying only the “steps/jumps” in the right-continuous step function that constitutes the ironed virtual valuation. That is, we implicitly extend the definition of each such ϕ_i over all of $[0, H]$ by defining $\phi_i(v_i) \triangleq \text{Max}\{\phi_i(w_i) \mid w_i \leq v_i \text{ \& } \phi_i(w_i) \text{ is defined}\}$ for every v_i , where this maximum is defined to be $-\infty$ (or any sufficiently small number) if it is taken over the empty set (i.e., this is the level of the leftmost “plateau” of ϕ_i). In particular, we note that the minimal winning bid of any bidder i in the Myersonian auction defined in Theorem 2.7 is always in $\text{supp } \hat{F}_i$.

3 Rounding Arbitrary Auctions

In this section, we derive Lemma 1.2(a) from Section 1. In fact, we prove a stronger result for arbitrary (not necessarily optimal) auctions. Fix an arbitrary auction A , fix $\varepsilon > 0$, and fix a product distribution over valuation profiles $F_1 \times \cdots \times F_n \in \Delta([0, H])^n$. In this section, we show that there exists an ε -rounding of A whose revenue from $F_1 \times \cdots \times F_n$ is less than an additive ε smaller than that of A .

We note that in Example A.1 in Appendix A, we show that the desired ε -rounding A' of A has to be constructed specifically for the target distribution $F_1 \times \cdots \times F_n$, i.e., that in some scenarios, every ε -rounding of A must lose $\Omega(1)$ (i.e., at least some constant independent of ε) revenue for some product distribution. We show this by showing that in such scenarios, every ε -rounding of A must lose $\Omega(1)$ revenue for some valuation profile. We show this even for the case where A is restricted to be the optimal auction for some $F_1 \times \cdots \times F_n$, even when F_1, \dots, F_n are regular⁶ distributions, even when the “rounded” auction A' may be any ε -coarse auction (whether or not an ε -rounding of A), and already

⁶Regularity is a standard tail assumption on distributions over valuations, which simplifies many analyses. We do not elaborate on its definition, as we do not require it for our analysis beyond remarking that all distributions in Example A.1 are regular.

for $n=2$ bidders. In contrast, in Appendix C we show that for the special case in which A is restricted to be the optimal auction for some i.i.d. product distribution $F^n = F \times \dots \times F$, it is always possible to construct an ε -rounding of A that loses less than an additive ε over any valuation profile, and thus over any distribution, by simply rounding-down the parameters representing the optimal auction.⁷ As shown in Appendix C, this implies a simplified proof for the recent result of Roughgarden and Schrijvers (2016) for the i.i.d. case, and highlights the difference between the general case considered here and the i.i.d. case.

Our proof strategy for Lemma 1.2(a) is a probabilistic one: Given an auction A and a target distribution $F = F_1 \times \dots \times F_n$, we will construct a distribution over ε -roundings of A that *in expectation* approximates the revenue of A from F , and hence deduce that at least one (deterministic) auction in the support of this distribution does not lose much revenue compared to A . We begin by defining this distribution over auctions, which rounds the bids in each ε -interval to a value that is randomly picked according to the restriction of the target distribution to that ε -interval.

Definition 3.1 (Conditional Distribution on an ε -Interval). For a distribution $F \in \Delta([0, H])$ and an index j of an ε -interval, we denote the conditional distribution of $v \sim F$ restricted to the ε -interval $[j \cdot \varepsilon, (j+1) \cdot \varepsilon)$ by $F|_j$.⁸

Definition 3.2 (Randomized ε -Rounding). Let $\varepsilon > 0$.

- Let $F_i \in \Delta([0, H])$. The F_i -randomized ε -rounding action on bidder i 's ε -interval j is a distribution over ε -rounding actions (i, j, v_j^i) obtained by randomly choosing $v_j^i \in [j \cdot \varepsilon, (j+1) \cdot \varepsilon)$ according to the conditional distribution $F_i|_j$ of F_i restricted to this ε -interval.⁹ Note that applying the F_i -randomized ε -rounding action to an auction A yields a distribution over deterministic auctions A' .
- Let $F = F_1 \times \dots \times F_n \in \Delta([0, H])^n$. The F -randomized ε -rounding rule is given by the collection of F_i -randomized ε -rounding actions for all bidders i and all ε -intervals j . Note that applying the F -randomized ε -rounding action to an auction A yields a distribution over deterministic ε -coarse auctions A' . This A' is called the F -randomized ε -rounding of A .

We start by analyzing the impact on the revenue from each bidder following the application of a single F_i -randomized ε -rounding action to some given auction. This is done in the following lemma, whose proof, once this lemma is stated, is quite straightforward (but is nonetheless spelled out in Appendix B, to which all omitted proofs are relegated).

Lemma 3.3. Let $F_1 \times \dots \times F_n \in \Delta([0, H])^n$, let A be an n -bidder auction, let $\varepsilon > 0$, and let $i \in N$. Let j be an index of an ε -interval, and A' be the result of applying the F_i -randomized ε -rounding action on bidder i 's interval j to A . Let $v_{-i} \in [0, H]^{N \setminus \{i\}}$ be a profile of valuations for all bidders other than i .

⁷The construction underlying Example A.1 in Appendix A also demonstrates that for non-i.i.d. product distributions, simply rounding-down all bids and applying the allocation rule of the optimal auction to the rounded-down bids (while adapting the payments to ensure truthfulness, of course) may lose $\Omega(1)$ in overall revenue. See Remark A.2 in Appendix A for more details.

⁸If this conditional distribution is ill defined, i.e., if $F([j \cdot \varepsilon, (j+1) \cdot \varepsilon)) = 0$, then we say that $F|_j$ is not defined.

⁹If $F_i|_j$ is not defined, then we pick $v_j^i \in [j \cdot \varepsilon, (j+1) \cdot \varepsilon)$ arbitrarily, say, $v_j^i \triangleq j \cdot \varepsilon$.

a. Let $i' \in N \setminus \{i\}$.

i. For every $v_i \notin [j \cdot \varepsilon, (j+1) \cdot \varepsilon)$, is it is surely¹⁰ the case that

$$r_{i'}^{A'}(v_1, \dots, v_n) = r_{i'}^A(v_1, \dots, v_n).$$

ii. If $F_i|_j$ is defined, then

$$\mathbb{E}_{v_i \sim F_i|_j} \left[\mathbb{E}_{A'} r_{i'}^{A'}(v_1, \dots, v_n) \right] = \mathbb{E}_{v_i \sim F_i|_j} \left[r_{i'}^A(v_1, \dots, v_n) \right].$$

b. Let w_i be the minimal winning bid of i in A when all other bidders bid v_{-i} .¹¹

i. If $w_i \notin (j \cdot \varepsilon, (j+1) \cdot \varepsilon)$,¹² then for every $v_i \in [0, H]$, it is surely the case that

$$r_i^{A'}(v_1, \dots, v_n) = r_i^A(v_1, \dots, v_n).$$

ii. Assume that $w_i \in (j \cdot \varepsilon, (j+1) \cdot \varepsilon)$.¹²

1. For every $v_i \notin [j \cdot \varepsilon, (j+1) \cdot \varepsilon)$, it is surely the case that v_i wins against v_{-i} in A' if and only if v_i wins against v_{-i} in A .
2. If $F_i|_j$ is defined, then the winning probability of a bid $v_i \sim F_i|_j$ against v_{-i} is the same in A' and in A (where in the former, the probability is taken over the randomness of both v_i and A' , and in the latter — over the randomness of v_i).
3. Let w'_i be the minimal winning bid of i in A' when all other bidders bid v_{-i} (so w'_i is a random variable). It is surely the case that

$$|w'_i - w_i| < \varepsilon.$$

Lemma 3.3 implies the following theorem, which we believe to be of independent interest.

Theorem 3.4. For every $F = F_1 \times \dots \times F_n \in \Delta([0, H])^n$, for every n -bidder auction A , and for every $\varepsilon > 0$, it is the case that

$$\left| \mathbb{E}_{A'} \text{Rev}^{A'}(F) - \text{Rev}^A(F) \right| < p \cdot \varepsilon,$$

where A' is the F -randomized ε -rounding of A , and where p is the probability that some bidder wins in A when the profile of bids is distributed according to F .¹³

¹⁰I.e., for every realization in the support of A' .

¹¹The fact that A is a direct-revelation DSIC and ex-post IR auction implies that a winning bidder pays her minimal winning bid, which is determined by v_{-i} .

¹²This is indeed an *open* interval. There is no typo here.

¹³In fact, a slightly stronger statement also holds, where p is replaced with the probability that some bidder wins in A and pays a price that is not an integer multiple of ε when the profile of bids is distributed according to F . See the proof of Theorem 3.4 for details. This implies, in a sense, that the revenue loss due to ε -rounding is smaller if A already behaves similarly to an ε -coarse auction for a set of valuation profiles that has positive probability.

It is useful, though, not to merely consider Theorem 3.4 as a consequence of Lemma 3.3, but to consider “intermediate results” between these two, both for instructive purposes (indeed, the deduction of Theorem 3.4 from Lemma 3.3 involves quite a few summations), but also as these results are of independent value. For example, Theorem 3.4 may be obtained by Lemma 3.3 by first considering the overall effect of a single randomized rounding action on the overall expected revenue.

Lemma 3.5. *Let $F = F_1 \times \cdots \times F_n \in \Delta([0, H])^n$, let A be an n -bidder auction, let $\varepsilon > 0$, let $i \in N$, and let j be an index of an ε -interval. Then,*

$$\left| \mathbb{E}_{A'} \text{Rev}^{A'} - \text{Rev}^A \right| < \varepsilon \cdot p_j^i,$$

where A' is the result of applying the F_i -randomized ε -rounding action on bidder i 's ε -interval j to A , and where p_j^i is the probability that i both wins and pays a price in $(j \cdot \varepsilon, (j + 1) \cdot \varepsilon)$ in A when the profile of bids is distributed according to F .

Lemma 3.5 will also be useful in Section 4. Another way to obtain Theorem 3.4 from Lemma 3.3 is to first consider the impact of the randomized rounding rule (i.e., all randomized rounding actions together), albeit only on the revenue from a single bidder i and only when her bid is restricted to some ε -interval.

Lemma 3.6. *Let $F = F_1 \times \cdots \times F_n \in \Delta([0, H])^n$, let A be an n -bidder auction, and let $\varepsilon > 0$. Let A' be the F -randomized ε -rounding of A . For every $i \in N$ and for every index j of an ε -interval such that $F_i|_j$ is defined, it is the case that*

$$\left| \mathbb{E}_{v_i \sim F_i|_j} \left[\mathbb{E}_{A'} r_i^{A'}(v_i) - r_i^A(v_i) \right] \right| < \varepsilon \cdot p,$$

where p is the probability that i wins in A when bidding $v_i \sim F_i|_j$, when the remaining bids are distributed according to F_{-i} .¹⁴

Recall that we show in Example A.1 that an ε -coarse, deterministic valuation-by-valuation approximation for A cannot exist for a product of general (even regular) non-i.i.d. distributions. Lemma 3.6 precisely reveals the gap between the approximation developed in this section and this (unattainable) goal: The approximation developed in this section provides a randomized (rather than deterministic) approximation that is amortized over each ε -interval of valuations (rather than attained on each valuation profile separately). I.e., the difference between the revenues of A and of its F -randomized ε -rounding from any bidder i , in expectation over $v_i \sim F_i|_j$ for any index j of an ε -interval, is small.

Finally, by Theorem 3.4, at least one deterministic auction in the support of A' loses less than an additive $p \cdot \varepsilon \leq \varepsilon$ in (overall) revenue compared to A , as formalized in Proposition 3.7 below, which is a restatement of Lemma 1.2(a).

Proposition 3.7 (Restatement of Lemma 1.2(a)). *For every product distribution $F = F_1 \times \cdots \times F_n \in \Delta([0, H])^n$ and for every auction A , there exists a (deterministic, ε -coarse) ε -rounding A' of A whose revenue from F is less than an additive ε smaller than that of A .*

¹⁴In fact, a slightly stronger statement also holds, where p is replaced with the probability that i wins in A and pays a price that is not an integer multiple of ε when bidding $v_i \sim F_i|_j$, when the remaining bids are distributed according to F_{-i} . See the proof of Lemma 3.6 for details.

Proof. Let A' be the F -randomized ε -rounding of A . By Theorem 3.4, we have that $\mathbb{E}_{A'} \text{Rev}^{A'}(F) > \text{Rev}^A(F) - \varepsilon$. Therefore, at least one of the deterministic (ε -coarse, ε -roundings of A) realizations of A' has revenue greater than $\text{Rev}^A(F) - \varepsilon$. \square

We note that in Example A.3 in Appendix A, we show that the bound of “less than ε ” on the revenue difference in Lemma 1.2(a) (and hence also in Proposition 3.7, in Theorem 3.4, etc.) cannot be unconditionally tightened any further, even when A' may be any ε -coarse auction (whether or not an ε -rounding of A), and already for $n=1$ bidder.

4 Efficiently Rounding Myersonian Auctions for Empirical Distributions

In the previous section, we have seen that for any target distribution $F = F_1 \times \dots \times F_n$, the revenue-maximizing auction A can be ε -rounded into an ε -coarse auction A' that approximately maximizes the revenue from F . In this section, we derive Lemma 1.2(b) from Section 1, showing how to deterministically and efficiently find a “good enough” ε -rounding of A when for every $i \in N$, it is the case that $F_i = \hat{F}_i$ is a uniform distribution over a multiset of t (sampled) values. (Recall that in such a scenario, the revenue-maximizing auction A can be computed in polynomial time using Theorem 2.7.) First, we note that if A is a Myersonian auction, then any rounding of A is also Myersonian, and its ironed virtual valuations can be efficiently computed from those of A , given the rounding actions.

Remark 4.1. Let $A = (\phi_i)_{i \in N}$ be an n -bidder Myersonian auction and let $\varepsilon > 0$.

- Applying an ε -rounding action (i, j, v_j^i) to A yields the Myersonian auction A' obtained by modifying ϕ_i to have the constant value $\phi_i(v_j^i)$ on the ε -interval $[j \cdot \varepsilon, (j+1) \cdot \varepsilon)$.
- Applying an ε -rounding rule to A (recall that an ε -rounding rule is a collection of ε -rounding actions (i, j, v_j^i) , one action for each pair of bidder i and index j of an ε -interval) yields the Myersonian auction A' obtained by modifying, for each bidder $i \in N$, the function ϕ_i to be constant on each ε -interval, having the constant value $\phi_i(v_j^i)$ on $[j \cdot \varepsilon, (j+1) \cdot \varepsilon)$ for each j .
- In particular, any ε -rounding of any Myersonian auction is also Myersonian (and ε -coarse).

So, how do we find a “good enough” rounding of the optimal Myersonian auction A ? One option, which would lead to a *randomized* polynomial-time algorithm for finding a suitable rounding with high probability, is to simply draw a polynomial number of deterministic roundings from the F -randomized ε -rounding of A , and pick the one that gives highest approximate (empirical) revenue among all drawn roundings, where the (empirical) revenue from a given rounding can be estimated by drawing a polynomial number of valuation profiles from F and averaging the (empirical) revenue from the given rounding over them. We next show that when for every $i \in N$, it is the case that $F_i = \hat{F}_i$ is a uniform distribution over a multiset of t (sampled) values, this process

can be derandomized to yield an efficient deterministic algorithm that finds a suitable rounding.¹⁵

We start with the question of deterministically computing the (empirical) revenue $\text{Rev}^{A'}(\hat{F})$ of a given auction A' , where $\hat{F} = \hat{F}_1 \times \cdots \times \hat{F}_n$. We note that in Example A.4 in Appendix A, we show that for general (even ε -coarse) auctions A' , this is a computationally hard problem, even when for every i it is the case that $|\text{supp } \hat{F}_i| = 2$. (This indeed motivates restricting our attention to Myersonian auctions in this section.) Nonetheless, we now show that for a Myersonian auction A' , its empirical revenue may be precisely computed in polynomial time. The key idea behind this result is that the revenue of a Myersonian auction from a given valuation profile only depends on the two highest ironed virtual bids; therefore, the algorithm efficiently calculates the overall revenue by enumerating over all of the possibilities for these two bids.

Lemma 4.2. *Let $n, t \in \mathbb{N}$. There exists a deterministic algorithm that runs in time $\text{poly}(n, t)$, such that given n discrete distributions $\hat{F}_1, \dots, \hat{F}_n$, each with support of size at most t , and for each $i \in N$ given a nondecreasing $\phi_i : \text{supp } \hat{F}_i \rightarrow \mathbb{R}$ (so, ϕ_i is a nondecreasing sequence of at most t real numbers), outputs the revenue $\text{Rev}^{(\phi_i)_{i \in N}}(\hat{F})$ of the Myersonian auction $(\phi_i)_{i \in N}$ over the product distribution $\hat{F} = \hat{F}_1 \times \cdots \times \hat{F}_n$.*

Proof. The algorithm is as follows:

1. For every $i \in N$, for every $j \in N \setminus \{i\}$, and for every $v_j \in \text{supp } \hat{F}_j$:
 - (a) If $j < i$, then let $L_j^i(v_j) = \mathbb{P}_{v_i \sim \hat{F}_i}(\phi_j(v_j) \geq \phi_i(v_i))$,
 - (b) If $j > i$, then let $L_j^i(v_j) = \mathbb{P}_{v_i \sim \hat{F}_i}(\phi_j(v_j) > \phi_i(v_i))$.

($L_j^i(v_j)$ is the probability that the value v_j of j precedes the value of i .)
2. Initialize: $r \leftarrow 0$.
3. For every $i \in N$, for every $j \in N \setminus \{i\}$, and for every $v_j \in \text{supp } \hat{F}_j$:
 - (a) If $i < j$, then let $w_i \in \hat{F}_i$ be minimal such that $\phi_i(w_i) \geq \phi_j(v_j)$,
 - (b) If $i > j$, then let $w_i \in \hat{F}_i$ be minimal such that $\phi_i(w_i) > \phi_j(v_j)$.
 - (c) Update: $r \leftarrow r + w_i \cdot (1 - L_j^i(v_j)) \cdot \hat{f}_j(v_j) \cdot \prod_{k \in N \setminus \{i, j\}} L_j^k(v_j)$.

(The added amount is the expected revenue from valuation profiles where bidder i wins and bidder j is the second-highest bidder with bid v_j .)
4. Output: r . □

We are now ready to present a deterministic algorithm that, given a Myersonian auction $A = (\phi_i)_{i \in N}$ that maximizes (empirical) revenue, finds an ε -rounding of A that maximizes (empirical) revenue up to less than an additive $n\varepsilon$. This algorithm sequentially and greedily searches for the values v_j^i for ε -rounding A , calculating the interim revenues (following each rounding action) via the algorithm of Lemma 4.2, and bounding the revenue loss via Lemma 3.5. From this algorithm, Lemma 1.2(b) follows.

¹⁵Another, in some sense less elegant, approach to derandomizing this algorithm would be to carefully use the randomness inherent in the order of the t given samples. We indeed use such an approach later when analyzing more general single-parameter auction environments in Section 6.3.

Proposition 4.3. *Let $n, t \in \mathbb{N}$ and let $\varepsilon > 0$. There exists a deterministic algorithm that runs in time $\text{poly}(H, n, 1/\varepsilon, t)$, such that given n discrete distributions $\hat{F}_1, \dots, \hat{F}_n$, each with support of size at most t , and for each $i \in N$ given a nondecreasing $\phi_i : \text{supp } \hat{F}_i \rightarrow \mathbb{R}$ (so, ϕ_i is a nondecreasing sequence of at most t real numbers), outputs for every $i \in N$ a function ϕ'_i , such that $(\phi'_i)_{i \in N}$ is an ε -rounding of $(\phi_i)_{i \in N}$, and such that*

$$\text{Rev}^{(\phi'_i)_{i \in N}}(\hat{F}) > \text{Rev}^{(\phi_i)_{i \in N}}(\hat{F}) - n\varepsilon,$$

where $\hat{F} = \hat{F}_1 \times \dots \times \hat{F}_n$.

Proof. The algorithm is as follows:

1. Initialize: $\phi'_i \leftarrow \phi_i$ for every i .
2. For every $j \in \{0, \dots, \lfloor H/\varepsilon \rfloor\}$: initialize¹⁶ $\phi'_i(j \cdot \varepsilon) \leftarrow \text{Max}_{w_i \in \text{supp } \hat{F}_i \cap [0, j \cdot \varepsilon]} \{\phi_i(w_i)\}$.
3. For every $i \in N$:
 - (a) For every $j \in \{0, \dots, \lfloor H/\varepsilon \rfloor\}$:
 - i. Let $\phi''_i \leftarrow \phi'_i$.
 - ii. For every $v_j^i \in \text{supp } \hat{F}_i \cap [j \cdot \varepsilon, (j+1) \cdot \varepsilon)$:
 - A. Let $\phi''_i(v_i) \leftarrow \phi'(v_j^i)$ for every $v_i \in (\text{supp } \hat{F}_i \cap [j \cdot \varepsilon, (j+1) \cdot \varepsilon)) \cup \{j \cdot \varepsilon\}$.
 - B. Compute $r_{v_j^i} \leftarrow \text{Rev}^{(\phi''_i, \phi'_{-i})}(\hat{F})$ via the algorithm of Lemma 4.2.
 - iii. For v_j^i for which $r_{v_j^i}$ is maximal, update $\phi'_i(v_i) \leftarrow \phi'(v_j^i)$ for every $v_i \in (\text{supp } \hat{F}_i \cap [j \cdot \varepsilon, (j+1) \cdot \varepsilon)) \cup \{j \cdot \varepsilon\}$.

(Lemma 3.5 guarantees that the revenue loss is less than $\varepsilon \cdot p_j^i$ for every pair (i, j) , and so less than $\varepsilon \cdot p_i \leq \varepsilon$ for every i , where p_i is the probability that i wins, for a total of less than $n\varepsilon$ over all bidders. Note that p_j^i does not change during our modification of ϕ_i at any ε -interval with index $k \neq j$; on the other hand, the probabilities p_i may well change during our modification of $\phi_{i'}$ for $i' \neq i$, and therefore our bound on the overall revenue loss is $n\varepsilon$ rather than ε as in the case of randomized rounding studied in Section 3.)

4. Output: $(\phi'_i)_{i \in N}$. □

Recall that we have shown in Section 3 that there exists an ε -rounding of any given Myersonian auction for any given target distribution, that loses less than an additive ε in revenue. It is unclear to us whether one can bridge the gap between this (tight) less-than- ε loss and the less-than- $n\varepsilon$ loss of Proposition 4.3 via an efficient deterministic algorithm.

5 Uniform Convergence over the Set of Rounded Myersonian Auctions

In this section, we conclude the derivation of Theorem 1.4 from Section 1. While most previous analyses (Morgenstern and Roughgarden, 2015; Roughgarden and Schrijvers,

¹⁶See Remark 2.8.

2016) restricted the set of possible “output auctions” to an infinite set and used advanced statistical tools to obtain that all auctions in this set perform similarly well on the empirical distribution and on the true distribution, in this paper we restrict the set of possible “output auctions” to a finite set, for which such a uniform-convergence result may be shown via an elegant concentration inequality due to Devanur et al. (2016).

Definition 5.1 (S_ε^n). For every $\varepsilon > 0$, we denote the set of all ε -coarse Myersonian n -bidder auctions (for valuations in $[0, H]$) by S_ε^n .

Lemma 5.2. $|S_\varepsilon^n| \leq \exp(\text{poly}(H, n, 1/\varepsilon))$.

*Proof.*¹⁷ Let $I \triangleq \{0, \dots, \lfloor H/\varepsilon \rfloor\}$. We claim that every ε -coarse Myersonian auction A' (for valuations in $[0, H]$) is completely specified by a sequence $((i_1, j_1), (i_2, j_2), \dots, (i_K, j_K))$ of distinct pairs from $N \times I$ (so $K \leq |N| \cdot |I|$). The correspondence is as follows: Given such a sequence, the allocation rule of the corresponding auction is specified by the following algorithm:

1. If $v_{i_1} \geq j_1 \cdot \varepsilon$, then i_1 wins; else, continue.
2. If $v_{i_2} \geq j_2 \cdot \varepsilon$, then i_2 wins; else, continue.
- \vdots
- K . If $v_{i_K} \geq j_K \cdot \varepsilon$, then i_K wins; else, continue.

$K+1$. No bidder wins.

The winner (if any) pays her minimal winning bid.

Indeed, such a sequence may be constructed for any ε -coarse Myersonian auction $A' = (\phi'_i)_{i \in N}$ by taking all pairs (i, j) such that ϕ'_i attains nonnegative value on the ε -interval with index j , and sorting these pairs in decreasing order of this nonnegative value, breaking ties in favor of lower i , and for the same i in favor of higher j .¹⁸ \square

The exponential size of S_ε^n means that if we can show for any given single auction that the number of samples required for it to perform similarly well on the empirical distribution \hat{F} and on the true distribution F is polynomial in H, n , and $1/\varepsilon$, but has logarithmic dependence on the desired success probability δ (as in the Chernoff-Hoeffding Inequality), then a union bound over S_ε^n can show that a polynomial number of samples suffices to guarantee that with high probability all ε -coarse Myersonian auctions perform similarly well on the empirical distribution \hat{F} and on the true distribution F . Fortunately, an elegant concentration inequality due to Devanur et al. (2016) shows precisely this (for any given single auction). The following proposition is a special case of this concentration inequality. (See also Proposition C.10 in Appendix C for an analogous yet somewhat more delicate concentration inequality for i.i.d. distributions, which does not follow from the analysis of Devanur et al., 2016.)

¹⁷This proof is reminiscent of the way Devanur et al. (2016) bound the number of optimal auctions on a finite valuation space.

¹⁸This results in a sequence with the additional property that for every $k \in \{1, \dots, K\}$ and for every $j \in \{j_k+1, \dots, \lfloor H/\varepsilon \rfloor\}$, there exists $\ell < k$ such that $i_\ell = i_k$ and $j_\ell = j$. The above correspondence between ε -coarse Myersonian auctions and sequences of bidder-index pairs, when the sequences are restricted to have this additional property, is in fact one-to-one.

Proposition 5.3 (Follows from Devanur et al., 2016). *For every $\varepsilon > 0$ and $\delta > 0$, there exists $t = \text{poly}(H, n, 1/\varepsilon, \log 1/\delta)$ such that the following holds. Fix $F_1, \dots, F_n \in \Delta([0, H])$ and fix an auction A . Draw t samples from each F_i , and let \hat{F}_i be the empirical uniform distribution over the t samples from F_i . Then, with probability at least $1 - \delta$ it is the case that*

$$|\text{Rev}^A(\hat{F}) - \text{Rev}^A(F)| < \varepsilon,$$

where $F = F_1 \times \dots \times F_n$ and $\hat{F} = \hat{F}_1 \times \dots \times \hat{F}_n$.

For completeness, we provide in the appendix a short proof of Proposition 5.3 via the Chernoff-Hoeffding Inequality and a union bound.

Remark 5.4. When we draw t tuples $(v_1^1, \dots, v_n^1), \dots, (v_1^t, \dots, v_n^t)$ from F , the empirical average value of the revenue of A is concentrated around its expectation $\text{Rev}^A(F)$. This is not what Proposition 5.3 (and Devanur et al., 2016) show, since this shows that $\mathbb{E}_{j \sim U(\{1, 2, \dots, t\})} r^A(v_1^j, \dots, v_n^j)$ is close to $\text{Rev}^A(F)$, while Proposition 5.3 shows that $\text{Rev}^A(\hat{F}) = \mathbb{E}_{j_1 \sim U(\{1, \dots, t\}), \dots, j_n \sim U(\{1, \dots, t\})} r^A(v_1^{j_1}, \dots, v_n^{j_n})$ is close to $\text{Rev}^A(F)$.

As noted above, the logarithmic dependence of t on $1/\delta$ in Proposition 5.3, in conjunction with the exponential size of S_ε^n , allows us to use the union bound to obtain the required uniform-convergence result over S_ε^n .

Lemma 5.5. *For every $\varepsilon > 0$ and $\delta > 0$, there exists $t = \text{poly}(H, n, 1/\varepsilon, \log 1/\delta)$ such that the following holds. Fix $F_1, \dots, F_n \in \Delta([0, H])$, draw t samples from each F_i , and let \hat{F}_i be the empirical uniform distribution over the t samples from F_i . With probability at least $1 - \delta$, it is the case that*

$$|\text{Rev}^A(F) - \text{Rev}^A(\hat{F})| < \varepsilon$$

holds simultaneously for every $A \in S_\varepsilon^n \cup \{OPT(F)\}$, where $F = F_1 \times \dots \times F_n$ and $\hat{F} = \hat{F}_1 \times \dots \times \hat{F}_n$.

Combining Lemma 5.5, Proposition 4.3, and Theorem 2.7, we obtain the following theorem, providing a polynomial-time algorithm for learning an approximately optimal auction from samples from an arbitrary unknown bounded product distribution.

Theorem 5.6 (Detailed version of Theorem 1.4). *There exists $t = \text{poly}(H, n, 1/\varepsilon, \log 1/\delta)$ such that the following holds. Let F_1, \dots, F_n be arbitrary distributions on $[0, H]$. For every $i \in N$, draw t samples from F_i , and let \hat{F}_i be the empirical distribution over the t samples from F_i . Then, with probability at least $1 - \delta$, the optimal auction for $\hat{F} = \hat{F}_1 \times \dots \times \hat{F}_n$ (which can be deterministically computed in time $\text{poly}(t)$ via Theorem 2.7), when $\frac{\varepsilon}{n+2}$ -rounded via the deterministic polynomial-time algorithm of Proposition 4.3, approximates the maximum possible revenue from $F = F_1 \times \dots \times F_n$ up to less than an additive ε .*

Proof. Let t be as in Lemma 5.5, for $S_{\varepsilon/(n+2)}^n$. By Theorem 2.7 and Proposition 4.3, all that we have to show is that the $\frac{\varepsilon}{n+2}$ -rounded auction $A' \in S_{\varepsilon/(n+2)}^n$ that maximizes the revenue from \hat{F} up to less than an additive $n \cdot \frac{\varepsilon}{n+2}$ also maximizes the revenue from F up to less than an additive ε . By Lemma 5.5, by definition of A' , and by Lemma 5.5 again, we have that

$$\text{Rev}^{OPT(F)}(F) < \text{Rev}^{OPT(F)}(\hat{F}) + \frac{\varepsilon}{n+2} \leq$$

$$\begin{aligned}
&\leq \text{Rev}^{OPT(\hat{F})}(\hat{F}) + \frac{\varepsilon}{n+2} < \\
&< \text{Rev}^{A'}(\hat{F}) + (n+1) \cdot \frac{\varepsilon}{n+2} < \\
&< \text{Rev}^{A'}(F) + (n+2) \cdot \frac{\varepsilon}{n+2} = \\
&= \text{Rev}^{A'}(F) + \varepsilon,
\end{aligned}$$

as required. \square

6 Beyond Single-Item Auctions

In this section, we extend the single-item analysis of the previous sections to more general single-parameter auction environments, deriving Theorem 1.5 from Section 1, as well as its generalization for intractable single-parameter auction environments. A *single-parameter (auction) environment* is defined by a set of *possible outcomes* $X \subseteq [0, 1]^n$, where a possible outcome $(x_1, \dots, x_n) \in X$ indicates that bidder 1 wins a fraction $x_1 \in [0, 1]$, bidder 2 wins a fraction $x_2 \in [0, 1]$, etc. The maximum willingness of a bidder i with valuation v_i to pay for a fraction x_i is $x_i \cdot v_i$. The single-item environment is obtained as a special case of single-parameter environment with

$$X_{\text{single-item}} = \{(0, \dots, 0), (1, 0, \dots, 0), (0, 1, 0, \dots, 0), \dots, (0, \dots, 0, 1)\}.$$

Additional important examples of single-parameter environments include *matroid environments* (see, e.g. Hartline, 2016), where the possible sets of winners correspond to independent sets in some matroid (N, I) over the ground set N of all bidders:

$$X_{(N, I)} = \{(x_1, \dots, x_n) \in \{0, 1\}^n \mid \{i \in N \mid x_i = 1\} \in I\};$$

and the (*nonexcludable*) *public project environment* (see, e.g. Hartline, 2016), where the auctioneer chooses whether or not to produce a public project from which all bidders benefit:

$$X_{\text{public-project}} = \{(0, \dots, 0), (1, \dots, 1)\}.$$

All of the above examples are special cases of *deterministic environments*, in which $X \subseteq \{0, 1\}^n$, i.e., each bidder either wins a fraction 1 or loses. An important example of a single-parameter environment that is not deterministic is that of *position environments* (see, e.g., Hartline, 2016; see also *sponsored search auctions* in Roughgarden, 2016), where n numbers $x^{(1)} \geq \dots \geq x^{(n)} \in [0, 1]$ are given, and each possible outcome corresponds to an ordering of the bidders, where the i th bidder in this order wins a fraction $x^{(i)}$.

6.1 Optimal Auctions

We once again make very weak use of Myerson's (1981) characterization of optimal auctions for single-parameter environments, and once again only present this characterization, which generalizes that presented in Section 2, to the extent required by our analysis.

Definition 6.1 (Myersonian Auction for a Single-Parameter Environment, Myerson, 1981). An n -bidder *single-parameter Myersonian auction* $((\phi_i)_{i \in N}; \text{MAX}_X)$ for the single-parameter environment $X \subseteq [0, 1]^n$ (and valuations in $[0, H]$) has, for every $i \in N$, a nondecreasing function $\phi_i : [0, H] \rightarrow \mathbb{R}$ called the *ironed virtual valuation* of bidder i . In this auction, the chosen outcome is $(x_1, \dots, x_n) \in X$ that maximizes the *ironed virtual*

welfare $\sum_{i=1}^n x_i \cdot \phi_i(v_i)$, where ties between maximizing outcomes are broken lexicographically;¹⁹ the payments are the unique²⁰ nonnegative payments that make this auction truthful for all valuation profiles in $[0, H]^n$.

For arbitrary single-parameter environments, Myerson (1981) proves the following generalization of Theorem 2.6.

Theorem 6.2 (Myerson, 1981). *For every single-parameter environment $X \subseteq [0, 1]^n$ and product distribution $F = F_1 \times \dots \times F_n$, the Myersonian auction $\text{OPT}(F; X) = ((\phi_i)_{i \in N}; \text{MAX}_X)$, where for every $i \in N$ the ironed virtual valuation ϕ_i is as in Theorem 2.6 (and so depends only on F_i , and can be efficiently computed for discrete F_i as in Theorem 2.7), achieves maximal revenue from F among all possible auctions for X .*

The reader may verify that the Myersonian auction $((\phi_i)_{i \in N}; \text{MAX}_{X_{\text{single-item}}})$ coincides with the Myersonian auction $(\phi_i)_{i \in N}$ defined in Definition 2.5.

6.2 Rounding Arbitrary Auctions

Analysis similar to that of Section 3 can be used to prove the following generalization of Theorem 3.4.

Theorem 6.3. *For every $F = F_1 \times \dots \times F_n \in \Delta([0, H])^n$, for every n -bidder auction A for a single-parameter environment $X \subseteq [0, 1]^n$, and for every $\varepsilon > 0$, it is the case that*

$$\left| \mathbb{E}_{A'} \text{Rev}^{A'}(F) - \text{Rev}^A(F) \right| < W_A \cdot \varepsilon,$$

where A' is the F -randomized ε -rounding of A , and where $W_A \triangleq \mathbb{E}_{(x_i)} \sum_{i=1}^n x_i \leq n$, where the expectation is over the outcome $(x_i)_{i \in N} \in X$ of A when the profile of bids is distributed according to F . (E.g., for a deterministic environment $X \subseteq \{0, 1\}^n$, W_A is the expected number of winners in the outcome of A .)²¹

By Theorem 6.3, we obtain the following generalization of Proposition 3.7.

Proposition 6.4. *For every product distribution $F = F_1 \times \dots \times F_n \in \Delta([0, H])^n$ and for every auction A for a single-parameter environment $X \subseteq [0, 1]^n$, there exists a (deterministic, ε -coarse) ε -rounding A' of A whose revenue from F is less than an additive $W_X \cdot \varepsilon$ smaller than that of A , where $W_X \triangleq \text{Max}_{(x_i) \in X} \sum_{i=1}^n x_i \leq n$. (E.g., for a deterministic environment $X \subseteq \{0, 1\}^n$, W_X is the maximal number of winners in any outcome in X .)*

6.3 Efficiently Rounding Myersonian Auctions for Empirical Distributions

While the deterministic rounding algorithm of Proposition 4.3 can be adapted to certain single-parameter environments beyond the single-item environment, such as matroid and

¹⁹Once again, we use lexicographic tie breaking for simplicity.

²⁰For details, see, e.g., Hartline (2016); Roughgarden (2016).

²¹In fact, as is the case for Theorem 3.4, a slightly stronger statement also holds, where W_A is replaced with, roughly speaking, the expectation over the sum of the fractions of the winnings whose threshold prices are not integer multiples of ε . This implies here as well, in a similar sense, that the revenue loss due to ε -rounding is smaller if A already behaves similarly to an ε -coarse auction for a set of valuation profiles that has positive probability.

position environments, its complexity even for these two environments becomes exponential in the maximal number of winners. For this reason, we now consider a different approach for efficiently searching for a revenue-approximating deterministic rounding of a given Myersonian auction A , for a given target distribution F . Recall that in Section 4, we sketched the following outline for a simple *randomized* polynomial-time algorithm for finding, with high probability, a revenue-approximating ε -rounding of A : draw a polynomial number of deterministic roundings from the F -randomized ε -rounding of A , and pick the one that gives the highest approximate revenue among all drawn roundings, where the revenue from a given rounding can be estimated by drawing a polynomial number of valuation profiles from F and averaging the revenue from the given rounding over them. The approach that we now present uses precisely this algorithm, however deterministically obtains the necessary randomness from a polynomial number of samples that are drawn from the true distribution of valuations, thus succeeding with high probability over the drawn samples. Since our entire deterministic algorithm from Theorem 5.6/Theorem 1.4 is only guaranteed to succeed with high probability over samples drawn from the true distribution of valuations, therefore utilizing a randomized algorithm to choose a suitable rounding of the empirical revenue-maximizing auction A in the above-described way (deterministically obtaining the necessary randomness from samples drawn from the true distribution) does not qualitatively change the nature of our overall result: a deterministic algorithm that runs in polynomial time and, given polynomially many samples from the true distribution, outputs a (deterministic) auction that approximately maximizes revenue with high probability, where the high probability is over the randomness of the given samples.

We start by formalizing the above randomized algorithm, to obtain an analogue of Proposition 4.3, which on one hand generalizes Proposition 4.3 beyond single-item auctions, but on the other hand trades determinism for randomization.

Definition 6.5 (Tractable Single-Parameter Environment). We say that a single-parameter environment $X \subseteq [0, 1]^n$ is *tractable* if there exists a deterministic algorithm that runs in time $\text{poly}(H, n)$, such that given the (realized) ironed virtual values $\phi_i(v_i)$ of all bidders, outputs the (lexicographically minimal) outcome $(x_1, \dots, x_n) \in X$ that maximizes the ironed virtual welfare $\sum_{i=1}^n x_i \cdot \phi_i(v_i)$.

Lemma 6.6. Fix a tractable single-parameter environment $X \subseteq [0, 1]^n$. Let $n, t \in \mathbb{N}$ and let $\varepsilon > 0$ and $\delta > 0$. There exists a randomized algorithm that runs in time $\text{poly}(H, n, 1/\varepsilon, \log 1/\delta, t)$ and uses at most $b = \text{poly}(H, n, 1/\varepsilon, \log 1/\delta, \log t)$ random bits, such that given n discrete distributions $\hat{F}_1, \dots, \hat{F}_n$, each with support of size at most t , and for each $i \in N$ given a nondecreasing $\phi_i : \text{supp } \hat{F}_i \rightarrow \mathbb{R}$ (so, ϕ_i is a nondecreasing sequence of at most t real numbers), outputs for every $i \in N$ a function $\phi'_i : \text{supp } \hat{F}_i \rightarrow \mathbb{R}$, such that $((\phi'_i)_{i \in N}; \text{MAX}_X)$ is an ε -rounding of $((\phi_i)_{i \in N}; \text{MAX}_X)$, and such that with probability at least $1 - \delta$ it is the case that

$$\text{Rev}^{((\phi'_i)_{i \in N}; \text{MAX}_X)}(\hat{F}) > \text{Rev}^{((\phi_i)_{i \in N}; \text{MAX}_X)}(\hat{F}) - (W_X + 1) \cdot \varepsilon,$$

where $\hat{F} = \hat{F}_1 \times \dots \times \hat{F}_n$.

While for some (true) distributions F (with very low entropy), it is not possible to extract $b = \text{poly}(H, n, 1/\varepsilon, \log 1/\delta, \log t)$ random bits from polynomially many random samples from F , we now show that such distributions can be easily identified, and an approximately optimal auction can easily be directly learned for them. We are therefore

able to derandomize the algorithm from Lemma 6.6 and obtain the following counterpart for Proposition 4.3 for general tractable single-parameter environments.

Proposition 6.7. *Fix a tractable single-parameter environment $X \subseteq [0, 1]^n$. Let $n, t \in \mathbb{N}$, let $\varepsilon > 0$ and $\delta > 0$. There exists $s = \text{poly}(H, n, 1/\varepsilon, \log 1/\delta, t) \geq t$ and a deterministic algorithm that runs in time $\text{poly}(H, n, 1/\varepsilon, \log 1/\delta, t)$, such that the following holds. Let F_1, \dots, F_n be arbitrary distributions on $[0, H]$. For every $i \in N$, draw s samples from F_i , and let \hat{F}_i be the empirical distribution over the first t of the s samples from F_i . The algorithm, given the s samples drawn from $F = F_1 \times \dots \times F_n$, and for each $i \in N$ given a nondecreasing $\phi_i : \text{supp } \hat{F}_i \rightarrow \mathbb{R}$ (so, ϕ_i is a nondecreasing sequence of at most t real numbers), outputs a bit e and for every $i \in N$ a function ϕ'_i , where if $e = 1$ then $((\phi'_i)_{i \in N}; \text{MAX}_X)$ is an ε -rounding of $((\phi_i)_{i \in N}; \text{MAX}_X)$, such that with probability at least $1 - \delta$, one of the following holds:*

- $e = 1$ and $\text{Rev}^{((\phi'_i)_{i \in N}; \text{MAX}_X)}(\hat{F}) > \text{Rev}^{((\phi_i)_{i \in N}; \text{MAX}_X)}(\hat{F}) - (W_X + 1) \cdot \varepsilon$,
where $\hat{F} = \hat{F}_1 \times \dots \times \hat{F}_n$.
- $e = 0$ and $\text{Rev}^{((\phi'_i)_{i \in N}; \text{MAX}_X)}(F) > \text{Rev}^{\text{OPT}(F; X)}(F) - \varepsilon$.

6.4 Uniform Convergence over the Set of Rounded Myersonian Auctions

Recall that in the single-item environment, we obtained a uniform convergence result by noting that any ε -rounding of an optimal auction can be encoded by essentially only encoding the relative order of all ironed virtual values $\phi_i(v_i)$ for $i \in N$ and $v_i \in \text{supp } \hat{F}_i$, and also of 0. While such an encoding suffices for single-parameter environments where the outcome that maximizes ironed virtual welfare can be found via a greedy-by-ironed-virtual-value algorithm that at each step adds the next compatible bidder with maximum nonnegative ironed virtual value $\phi_i(v_i)$ to the set of winners, such as matroid and position environments, this encoding is unsuitable for more general environments. Indeed, already for the public project environment, knowledge of the order of all ironed virtual values $\phi_i(v_i)$ and 0 does not suffice in order to determine for a given valuation profile $(v_1, \dots, v_n) \in \text{supp } \hat{F}$ whether the ironed virtual welfare $\sum_{i=1}^n \phi_i(v_i)$ for the outcome where all bidders win is greater than or less than 0, the ironed virtual welfare for the outcome where no bidder wins. We now show that nonetheless, for any deterministic environment (including the public project environment), if we encode every ε -rounding of a Myersonian auction for this environment using the relative order of the ironed virtual welfares for all possible outcomes and all valuation profiles, then the set of all possible encodings still has size at most $\exp(\text{poly}(H, n, 1/\varepsilon))$, and so the analysis of Section 5 may be used to obtain the required an appropriate uniform convergence result for arbitrary deterministic environments.

Definition 6.8 (S_ε^X). For every $\varepsilon > 0$, we denote the set of all ε -coarse Myersonian n -bidder auctions for a given single-parameter environment $X \subseteq [0, 1]^n$ by S_ε^X .

Lemma 6.9. *There exists $e \leq \exp(\text{poly}(H, n, 1/\varepsilon))$ such that $|S_\varepsilon^X| \leq e$ for every single-parameter deterministic environment $X \subseteq \{0, 1\}^n$.*

²²This proof uses a technique in the spirit of the one popularly used to count weighted linear threshold circuits. We are not aware of the use of any similar argument in the literature on mechanism design, and hope that this type of argument may find additional uses in similar contexts in the future.

*Proof.*²² Let $((\phi'_i)_{i \in N}; \text{MAX}_X)$ be an ε -coarse Myersonian n -bidder auction for X . Let $M \triangleq n \cdot (\lfloor H/\varepsilon \rfloor + 1)$. We will show that for every $i \in N$, there exists a function $\psi'_i : [0, H] \rightarrow \mathbb{R}_+$ that is constant on each ε -interval and only attains values that are signed rational numbers with nominator and denominator both having absolute value at most $M!$, such that $((\psi'_i)_{i \in N}; \text{MAX}_X) = ((\phi'_i)_{i \in N}; \text{MAX}_X)$, i.e., these two Myersonian auctions have the same outcome for every valuation profile. As this implies that $|S_\varepsilon^X| \leq (2 \cdot M!^2 + 1)^M \leq \exp(\text{poly}(H, n, 1/\varepsilon))$, the lemma will follow.

For every $i \in N$ and $j \in \{0, \dots, \lfloor H/\varepsilon \rfloor\}$, let $\phi'_{i,j}$ be the value that ϕ'_i attains on the ε -interval $[j \cdot \varepsilon, (j+1) \cdot \varepsilon)$. We will define the functions ψ'_1, \dots, ψ'_n by defining the corresponding sequence of M values $\psi'_{i,j}$ via a set of linear constraints as follows: For every $j_1, \dots, j_n \in \{0, \dots, \lfloor H/\varepsilon \rfloor\}$ and for every two distinct subsets $N_1, N_2 \subseteq N$, we add the following constraint:

$$\begin{array}{lll} \sum_{i \in N_1} \psi'_{i,j_i} = \sum_{i \in N_2} \psi'_{i,j_i} & \text{if} & \sum_{i \in N_1} \phi'_{i,j_i} = \sum_{i \in N_2} \phi'_{i,j_i}, \\ \sum_{i \in N_1} \psi'_{i,j_i} \leq \sum_{i \in N_2} \psi'_{i,j_i} - 1 & \text{if} & \sum_{i \in N_1} \phi'_{i,j_i} < \sum_{i \in N_2} \phi'_{i,j_i}, \\ \sum_{i \in N_1} \psi'_{i,j_i} \geq \sum_{i \in N_2} \psi'_{i,j_i} + 1 & \text{if} & \sum_{i \in N_1} \phi'_{i,j_i} > \sum_{i \in N_2} \phi'_{i,j_i}. \end{array}$$

We note that the set of constraints that we have defined is satisfiable. Indeed, setting

$$d \triangleq \min \left\{ \left| \sum_{i \in N_1} \phi'_{i,j_i} - \sum_{i \in N_2} \phi'_{i,j_i} \right| \mid \begin{array}{l} j_1, \dots, j_n \in \{0, \dots, \lfloor H/\varepsilon \rfloor\} \\ \& N_1, N_2 \subseteq N \end{array} \& \sum_{i \in N_1} \phi'_{i,j_i} \neq \sum_{i \in N_2} \phi'_{i,j_i} \right\},$$

where by finiteness $d > 0$, we note that the assignment $\psi'_{i,j} = \frac{\phi'_{i,j}}{d}$ for all i, j satisfies this set of constraints. Since this set of (exponentially many) constraints in M unknowns is satisfiable, there exists a solution defined by precisely M of these constraints, where all M constraints are binding (i.e., hold with an equality rather than with an inequality). Let $(\psi'_{i,j})_{i,j}$ be this solution. By Cramer's Rule, each $\psi'_{i,j}$ is the quotient of the determinants of two $M \times M$ matrices, where in our case, by construction each of these matrices contains only values in $\{-1, 0, 1\}$, and therefore, each such determinant is an integer having absolute value at most $M!$. We conclude the proof as by construction, for every valuation profile, the order of the ironed virtual welfares of every two possible deterministic outcomes in $\{0, 1\}^n$ is the same according to both $(\phi'_i)_{i \in N}$ and $(\psi'_i)_{i \in N}$, and therefore $((\psi'_i)_{i \in N}; \text{MAX}_X) = ((\phi'_i)_{i \in N}; \text{MAX}_X)$, as required. \square

Plugging in Theorem 6.2, Propositions 6.4 and 6.7, and Lemma 6.9 into the analysis of the previous sections, we obtain the following analogue of Theorem 5.6, providing a polynomial-time algorithm for learning an approximately optimal auction for a class of single-parameter environments that includes deterministic environments (including matroid environments and the public project environment) and position environments, from samples from an arbitrary unknown bounded product distribution.

Theorem 6.10 (Detailed version of Theorem 1.5). *Let $X \subseteq [0, 1]^n$ be a tractable deterministic environment (such as a matroid environment or the public project environment) or a position environment. There exist $t \leq s = \text{poly}(H, n, 1/\varepsilon, \log 1/\delta)$ such that the following holds. Let F_1, \dots, F_n be arbitrary distributions on $[0, H]$. For every $i \in N$, draw s samples from F_i , and let \hat{F}_i be the empirical distribution over the first t of the s samples from F_i . Then, with probability at least $1 - \delta$, the auction output by the deterministic polynomial-time “ $\frac{\varepsilon}{W_X + 3}$ -rounding” algorithm of Proposition 6.7 given the s samples from*

$F = F_1 \times \cdots \times F_n$ and given the optimal auction for $\hat{F} = \hat{F}_1 \times \cdots \times \hat{F}_n$ (which by Theorem 6.2 can be deterministically computed in time $\text{poly}(t)$ via Theorem 2.7), approximates the maximum possible revenue from F up to less than an additive ε .

Remark 6.11. Theorem 6.10 also holds, via the same proof, for every tractable single-parameter environment $X \subseteq [0, 1]^n$ such that $|S_\varepsilon^X| \leq \exp(\text{poly}(H, n, 1/\varepsilon))$.

6.5 Computationally Hard Auction Environments

Up until now, we have assumed that the given single-parameter environment $X \subseteq [0, 1]^n$ is tractable. Nonetheless, for many single-parameter environments, including deterministic environments, this is known not to be the case. We therefore now extend our analysis to such environments.

A notable example of an intractable single-parameter environment is that of knapsack environments (Mu'alem and Nisan, 2008; see also Roughgarden, 2016), where each bidder i has a nonnegative “size” w_i , and a deterministic outcome is possible if the cumulative size of all winners does not exceed some threshold w :

$$X_{(w_1, \dots, w_n; w)} = \left\{ (x_1, \dots, x_n) \in \{0, 1\}^n \mid \sum_{i=1}^n w_i \cdot x_i \leq w \right\}.$$

In such environments, to precisely maximize ironed virtual welfare (and therefore revenue), one must solve an instance of the well known KNAPSACK problem, which is known to be NP-complete.

A popular approach in the literature for designing auctions for intractable environments X is to have the designed auction choose an outcome that approximately maximizes the ironed virtual welfare using some efficient constant-factor approximation algorithm APPROX_X instead of the intractable precise-maximization algorithm MAX_X ; the payments are once again the unique nonnegative payments that ensure (precise) truthfulness²³ for all valuation profiles in $[0, H]^n$. We denote this auction, which by tools developed by Myerson (1981) turns out to approximate the revenue from the optimal auction $\text{OPT}(F; X) = ((\phi_i)_{i \in N}; \text{MAX}_X)$ up to the same constant multiplicative factor of the approximation algorithm, by $((\phi_i)_{i \in N}; \text{APPROX}_X)$.

Assume, therefore, that $X \subseteq [0, 1]^n$ is an intractable environment and that APPROX_X is a deterministic algorithm that runs in time $\text{poly}(H, n)$, such that given the (realized) ironed virtual values $\phi_i(v_i)$ of all bidders, APPROX_X outputs an outcome $(x_1, \dots, x_n) \in X$ that maximizes the ironed virtual welfare $\sum_{i=1}^n x_i \cdot \phi_i(v_i)$ up to some multiplicative factor $C > 1$. We will explore when our analysis can be applied to yield a polynomial-time algorithm for learning, from samples from an arbitrary unknown bounded product distribution, a tractable auction that approximates the optimal revenue up to the same multiplicative factor of C , plus less than an additive ε .

We start by noting that the analysis and results of Section 6.2 hold for all auctions, and in particular also auctions of the form $A = ((\phi_i)_{i \in N}; \text{APPROX}_X)$. Furthermore, we note that the analysis of Section 6.3 still holds when replacing every occurrence of MAX_X with APPROX_X and every occurrence of $\text{Rev}^{\text{OPT}(F; X)}(F)$ with $\frac{\text{Rev}^{\text{OPT}(F; X)}(F)}{C}$.

²³For this auction to be truthful, the approximation algorithm must also maintain a certain well understood monotonicity property. For details, see, e.g., Hartline (2016); Roughgarden (2016).

Proposition 6.12. *Proposition 6.7 still holds when every occurrence of MAX_X is replaced with APPROX_X , when $\text{Rev}^{\text{OPT}(F;X)}(F)$ is replaced with $\frac{\text{Rev}^{\text{OPT}(F;X)}(F)}{C}$, and when requiring only that APPROX_X runs in time $\text{poly}(H, n)$ rather than requiring that X is tractable.*

Let $\text{OPT}(\hat{F}; X) = ((\phi_i)_{i \in N}; \text{MAX}_X)$ be the optimal auction for the empirical distribution $\hat{F} = \hat{F}_1 \times \dots \times \hat{F}_n$ in the environment X . By assumption, the truthful auction $A = ((\phi_i)_{i \in N}; \text{APPROX}_X)$ approximates the optimal revenue from \hat{F} up to a multiplicative factor of C . By Proposition 6.12, there exists an efficient algorithm that given $(\phi_i)_{i \in N}$ and \hat{F} (and given the ability to draw some additional samples from F , if the algorithm is to be deterministic rather than randomized), outputs an ε -rounding $A' = ((\phi'_i)_{i \in N}; \text{APPROX}_X)$ of A such that

$$\text{Rev}^{A'}(\hat{F}) > \text{Rev}^A(\hat{F}) - (W_X + 1) \cdot \varepsilon \geq \frac{\text{Rev}^{\text{OPT}(\hat{F}; X)}(\hat{F})}{C} - (W_X + 1) \cdot \varepsilon.$$

Now, if we had a result analogous to Lemma 5.5, showing that with high probability, for each ε -coarse auctions that uses APPROX_X to maximize some ironed virtual welfare as well as for $\text{OPT}(F; X)$, its revenues on F and on \hat{F} differ by less than an additive ε , then we would obtain the desired approximation result via a derivation similar to that of Theorem 6.10 (and Theorem 5.6):

$$\begin{aligned} \text{Rev}^{\text{OPT}(F; X)}(F) &< \text{Rev}^{\text{OPT}(F; X)}(\hat{F}) + \varepsilon \leq \\ &\leq \text{Rev}^{\text{OPT}(\hat{F}; X)}(\hat{F}) + \varepsilon < \\ &< C \cdot (\text{Rev}^{A'}(\hat{F}) + (W_X + 1) \cdot \varepsilon) + \varepsilon < \\ &< C \cdot (\text{Rev}^{A'}(F) + (W_X + 2) \cdot \varepsilon) + \varepsilon = \\ &= C \cdot (\text{Rev}^{A'}(F) + (W_X + 2 + 1/C) \cdot \varepsilon) \leq \\ &\leq C \cdot (\text{Rev}^{A'}(F) + (W_X + 3) \cdot \varepsilon). \end{aligned}$$

While a first glance may suggest that Lemma 6.9 indeed gives the desired analogue of Lemma 5.5, a closer look shows that more care is required here. Indeed, the fact that, as in the proof of Lemma 6.9, $((\psi'_i)_{i \in N}; \text{MAX}_X) = ((\phi'_i)_{i \in N}; \text{MAX}_X)$ by no means implies that $((\psi'_i)_{i \in N}; \text{APPROX}_X) = ((\phi'_i)_{i \in N}; \text{APPROX}_X)$, as many approximation algorithms use more information from $(\phi'_i)_{i \in N}$ beyond merely the order of all possible ironed virtual welfares for all possible valuation profiles and outcomes. For instance, the 2-approximation algorithm for the knapsack environment (Mu'alem and Nisan, 2008) also considers the order of the quotients $(\frac{\phi'_i(v_i)}{b_i})_{i \in N}$, which cannot be inferred from the order of all possible ironed virtual welfares alone. In the specific case of this 2-approximation algorithm, this is not the end of the road since the set of all possible orders of these quotients is of size at most $\exp(\text{poly}(H, n, 1/\varepsilon))$. More generally, we can obtain our approximation result whenever, for every ε -coarse $(\phi'_i)_{i \in N}$, we can encode all of the information that APPROX_X uses from $(\phi'_i)_{i \in N}$, such that the set of all possible encodings (i.e., the set of distinct ε -coarse auctions that use APPROX_X to maximize some ironed virtual welfare) is of size at most $\exp(\text{poly}(H, n, 1/\varepsilon))$. (This is the case, for example, if APPROX_X only reads the virtual valuations $\phi'_i(v_i)$ up to a precision of some polynomial number of bits.)

Definition 6.13 ($S_\varepsilon^{\text{APPROX}_X}$). For every $\varepsilon > 0$, we denote the set of all ε -coarse auctions of the form $((\phi_i)_{i \in N}; \text{APPROX}_X)$ for a given approximation algorithm APPROX_X for a given single-parameter environment $X \subseteq [0, 1]^n$ by $S_\varepsilon^{\text{APPROX}_X}$.

Theorem 6.14 (Generalization of Theorem 6.10/Theorem 1.5 for Intractable Auction Environments). *Let $X \subseteq [0, 1]^n$ be a single-parameter environment and let APPROX_X be a monotone algorithm that runs in time $\text{poly}(H, n)$ and finds an outcome in X that maximizes ironed virtual welfare up to some multiplicative factor $C \geq 1$. If $|S_\varepsilon^{\text{APPROX}_X}| \leq \exp(\text{poly}(H, n, 1/\varepsilon))$, then there exist $t \leq s = \text{poly}(H, n, 1/\varepsilon, \log 1/\delta)$ such that the following holds. Let F_1, \dots, F_n be arbitrary distributions on $[0, H]$. For every $i \in N$, draw s samples from F_i , and let \hat{F}_i be the empirical distribution over the first t of the s samples from F_i . Then, with probability at least $1 - \delta$, the auction output by the deterministic polynomial-time “ $\frac{\varepsilon}{W_X + 3}$ -rounding” algorithm of Proposition 6.12 given the s samples from $F = F_1 \times \dots \times F_n$ and given $((\phi_i)_{i \in N}; \text{APPROX}_X)$, for $(\phi_i)_{i \in N}$ that can be deterministically computed in time $\text{poly}(t)$ via Theorem 2.7 given $\hat{F} = \hat{F}_1 \times \dots \times \hat{F}_n$, approximates the maximum possible revenue from F up to the same multiplicative factor of C , plus less than an additive ε .*

We note that Theorem 6.14 is a strict generalization of Theorem 6.10. Indeed, fixing $C = 1$ in Theorem 6.14 yields Theorem 6.10, as a polynomial-time APPROX_X that guarantees revenue maximization “up to a multiplicative factor of 1” is precisely a tractable MAX_X .

7 Acknowledgments

Yannai Gonczarowski is supported by the Adams Fellowship Program of the Israel Academy of Sciences and Humanities. The work of Noam Nisan was supported by ISF grant 1435/14 administered by the Israeli Academy of Sciences and by Israel-USA Binational Science Foundation (BSF) grant number 2014389. We thank Jamie Morgenstern and Tim Roughgarden for explaining to us some of the subtleties of the problem of efficient empirical revenue maximization in single-item auctions.

References

- R. Cole and T. Roughgarden. The sample complexity of revenue maximization. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing (STOC)*, pages 243–252, 2014.
- N. R. Devanur, Z. Huang, and C.-A. Psomas. The sample complexity of auctions with side information. In *Proceedings of the 48th Annual ACM Symposium on Theory of Computing (STOC)*, pages 426–439, 2016.
- S. Dughmi, L. Han, and N. Nisan. Sampling and representation complexity of revenue maximization. In *Proceedings of the 10th Conference on Web and Internet Economics (WINE)*, pages 277–291, 2014.
- E. Elkind. Designing and learning optimal finite support auctions. In *Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 736–745, 2007.
- S. Hart and N. Nisan. The menu-size complexity of auctions. In *Proceedings of the 14th ACM Conference on Electronic Commerce (EC)*, page 565, 2013.

- J. D. Hartline. *Mechanism Design and Approximation*. 2016. Book draft.
- Z. Huang, Y. Mansour, and T. Roughgarden. Making the most of your samples. In *Proceedings of the 16th ACM Conference on Economics and Computation (EC)*, pages 45–60, 2015.
- J. H. Morgenstern and T. Roughgarden. On the pseudo-dimension of nearly optimal auctions. In *Advances in Neural Information Processing Systems 28 (NIPS)*, pages 136–144, 2015.
- A. Mu’alem and N. Nisan. Truthful approximation mechanisms for restricted combinatorial auctions. *Games and Economic Behavior*, 64(2):612–631, 2008.
- R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- T. Roughgarden. *Twenty Lectures on Algorithmic Game Theory*. Cambridge University Press, 2016.
- T. Roughgarden and O. Schrijvers. Ironing in the dark. In *Proceedings of the 17th ACM Conference on Economics and Computation (EC)*, pages 1–18, 2016.
- J. von Neumann. Various techniques used in connection with random digits. *National Bureau of Standards Applied Mathematics Series*, 12:36–38, 1951.

A Examples Omitted from the Main Text

Example A.1. Fix $H \triangleq 2$. For every $\varepsilon > 0$, there exists a product of regular distributions $F = F_1 \times F_2 \in \Delta([0, H])^2$ such that for every ε -coarse 2-bidder auction A' , there exists a valuation profile $(v_1, v_2) \in [0, H]^2$ such that $r^{A'}(v_1, v_2) < r^{OPT}(v_1, v_2) - 1/13$, where $OPT = OPT(F)$ is the Myersonian auction that maximizes the revenue from F .

Indeed, let $\varepsilon > 0$ and assume without loss of generality that $\varepsilon < 1/3$ (otherwise, divide ε by some large-enough integer). Let $\lfloor 1 \rfloor_\varepsilon$ denote the largest integer multiple of ε that is not greater than 1. Let $F_1 = U([0, 2])$ and let $F_2 \in \Delta([\lfloor 1 \rfloor_\varepsilon - \varepsilon, \lfloor 1 \rfloor_\varepsilon])$ be the continuous distribution with density function $f_2(x) = 2/\varepsilon^2 \cdot (x - \lfloor 1 \rfloor_\varepsilon + \varepsilon)$. By Theorem 2.6, $OPT = (\phi_1, \phi_2)$ for the (ironed) virtual valuations ϕ_1, ϕ_2 corresponding to F_1, F_2 , respectively. It is well known for the uniform distribution F_1 that the corresponding (ironed) virtual valuation satisfies $\phi_1(v_1) = 2 \cdot v_1 - 2$ for every $v_1 \in [0, 2]$. The crux of this example is that F_2 is defined so that despite the fact that it is tightly concentrated, the image of the corresponding (ironed) virtual valuation ϕ_2 is $(-\infty, \lfloor 1 \rfloor_\varepsilon]$. In particular, since $\varepsilon < 1/3$, there exist $v_2, w_2 \in [\lfloor 1 \rfloor_\varepsilon - \varepsilon, \lfloor 1 \rfloor_\varepsilon)$ such that $\phi_2(v_2) = 1/3$ and $\phi_2(w_2) = 2/3$. Note that $\phi_1(v_1) = 1/2$ for $v_1 = 5/4$, and $\phi_1(w_1) = 1$ for $w_1 = 3/2$. Therefore $r^{OPT}(v_1, v_2) = 7/6$ (bidder 1 wins and pays her minimal winning bid of $\phi_1^{-1}(1/3) = 7/6$) and $r^{OPT}(w_1, w_2) = 4/3$ (bidder 1 wins and pays her minimal winning bid $\phi_1^{-1}(2/3) = 4/3$).

Let A' be an ε -coarse 2-bidder auction. We will show that the proposition is satisfied for either (v_1, v_2) or (w_1, w_2) . Assume that $r^{A'}(v_1, v_2) \geq r^{OPT}(v_1, v_2) - 1/13$. Therefore, we have that $r^{A'}(v_1, v_2) > 1$. Since $v_2, w_2 \in [\lfloor 1 \rfloor_\varepsilon - \varepsilon, \lfloor 1 \rfloor_\varepsilon)$, we therefore have that $r^{A'}(v_1, w_2) = r^{A'}(v_1, v_2) > 1$ as well. Therefore, for the valuation profile (v_1, w_2) , the winner is bidder 1 (since bidder 2 never pays more than her value). So, when bidder 2

bids w_2 , the minimal winning bid of bidder 1 is at most $v_1 = 5/4$, and so bidder 1 does not pay more than $5/4$ if she wins against w_2 . Therefore,

$$r^{A'}(w_1, w_2) \leq 5/4 = 4/3 - 1/12 = r^{OPT}(w_1, w_2) - 1/12 < r^{OPT}(w_1, w_2) - 1/13,$$

as claimed.

Remark A.2. The construction underlying Example A.1 also demonstrates that for small enough $\varepsilon > 0$, there exists a product of regular distribution $F = F_1 \times F_2 \in \Delta([0, H])^2$ such that if A' is the auction obtained by simply rounding-down each bid to the nearest integer multiple of ε and applying the allocation rule of $OPT = OPT(F)$ to the rounded-down bids (while adapting the payments to ensure truthfulness),²⁴ then $\text{Rev}^{A'}(F) < \text{Rev}^{OPT}(F) - 1/8$.

Indeed, let $\varepsilon > 0$ and assume that $\varepsilon < 1/4$. For $F = F_1 \times F_2$ as in Example A.1, the “rounded-down” auction A' is such that bidder 1 wins if and only if $v_1 \geq \lceil 1 \rceil_\varepsilon$, where $\lceil 1 \rceil_\varepsilon$ is the smallest integer multiple of ε that is not less than 1, and bidder 2 almost surely²⁵ loses. Therefore, $\text{Rev}^{A'}(F) = \frac{2 - \lceil 1 \rceil_\varepsilon}{2} \cdot \lceil 1 \rceil_\varepsilon < 1/2 \cdot (1 + \varepsilon)$. We lower-bound the revenue from OPT by considering an auction that sells to bidder 1 if $v_1 \geq 3/2$, and otherwise to bidder 2. The revenue from this auction is $1/4 \cdot 3/2 + 3/4 \cdot (\lceil 1 \rceil_\varepsilon - \varepsilon) > 9/8 - 3/2 \cdot \varepsilon$. Therefore,

$$\text{Rev}^{OPT}(F) - \text{Rev}^{A'}(F) > 9/8 - 3/2 \cdot \varepsilon - 1/2 \cdot (1 + \varepsilon) = 5/8 - 2 \cdot \varepsilon > 1/8,$$

as claimed.

Example A.3. For every $H \geq \varepsilon > 0$ and $\varepsilon > \eta > 0$, there exists $F \in \Delta([0, H])$ such that for every ε -coarse 1-bidder auction A' , it is the case that $\text{Rev}^{A'}(F) \leq \text{Rev}^{OPT(F)}(F) - \eta$.

Indeed, let $v \triangleq \frac{\varepsilon + \eta}{2} > \eta > 0$ and let $p \triangleq \frac{\eta}{v} \in (0, 1)$. Let $F \in \Delta(\{0, v\})$ that assigns probability p to the value v and probability $1 - p$ to the value 0. Three truthful auctions (allocation rules) exist for F :

1. the bidder always wins (and pays her minimal winning bid of 0),
2. the bidder wins only if she bids v (and pays her minimal winning bid of v),²⁶ and
3. the bidder never wins.

The revenue (from F) of the first and third auctions is 0, while the revenue of the second auction is $p \cdot v = \eta > 0$. Therefore, $OPT(F)$ is the second of these three auctions. Note, however, that since $\{0, v\} \in [0, \varepsilon)$, only the first and third auctions are ε -coarse, so the revenue of any ε -coarse auction from F is 0, i.e., the loss in revenue of any ε -coarse auction, compared to $OPT(F)$, is an additive η (and in particular, at least an additive η), as claimed.

Example A.4. Let $F_1 = F_2 = \dots = F_n = U(\{0, 1\})$. For every $w_2, w_3, \dots, w_n \in \mathbb{N}$ and $w \in \mathbb{N}$, we define the n -bidder (DSIC and ex-post IR) auction $A'(w_2, w_3, \dots, w_n; w)$ to be the auction in which bidder 1 wins and pays 1 if and only if both $v_1 = 1$ and $\sum_{i=2}^n w_i \cdot v_i = w$ (otherwise, no bidder wins). Computing the revenue $\text{Rev}^{A'(w_2, \dots, w_n; w)}(F_1 \times \dots \times F_n)$ is #P-complete.

²⁴I.e., having each winning bidder pay her minimal winning bid.

²⁵I.e., with probability 1.

²⁶In fact, any payment in $[0, v]$ can be charged in this case, but for our analysis we will be interested in the auction of this form that achieves maximal revenue, and this auction charges v if the bidder wins.

We first note that $A'(w_2, \dots, w_n; w)$ is indeed DSIC and ex-post IR. Indeed, all bidders except bidder 1 have utility 0 regardless of their reported bids, and so have no incentive to misreport their values or to not participate. Furthermore, bidder 1 has utility 0 when bidding 0 and when truthfully bidding 1, and utility ≤ 0 when untruthfully bidding 1.

Now, by definition, the revenue $\text{Rev}^{A'(w_2, \dots, w_n; w)}(F_1 \times \dots \times F_n)$ equals half of the probability that $\sum_{i=2}^n w_i \cdot v_i = w$ for v_2, v_3, \dots, v_n that are drawn independently and uniformly from $\{0, 1\}$. This probability, in turn, equals exactly $2^{-(n-1)}$ times the number of subsets of $\{w_2, \dots, w_n\}$ whose sum equals w . To compute the latter, one must solve an instance of the counting analogue of the well known SUBSET-SUM problem, which is known to be $\#P$ -complete.

B Proofs Omitted from the Main Text

B.1 Proofs for Section 3

Proof of Lemma 3.3. Assume without loss of generality that $i = 1$. We prove each part separately.

Proof of Part a(i): When $v_1 \notin [j \cdot \varepsilon, (j+1) \cdot \varepsilon)$, the rounding of bids of bidder 1 in the ε -interval with index j affects neither the determination of the winner, nor the minimal winning bid (which equals the payment) of bidders other than bidder 1 (e.g., bidder i'), as required.

Proof of Part a(ii):

$$\mathbb{E}_{v_1 \sim F_1|j} \left[\mathbb{E}_{A'} r_{i'}^{A'}(v_1, \dots, v_n) \right] = \mathbb{E}_{v_j^1 \sim F_1|j} \left[r_{i'}^A(v_j^1, v_2, v_3, \dots, v_n) \right] = \mathbb{E}_{v_1 \sim F_1|j} \left[r_{i'}^A(v_1, v_2, \dots, v_n) \right].$$

Proof of Part b(i): Assume that $w_1 \notin (j \cdot \varepsilon, (j+1) \cdot \varepsilon)$. Therefore, the minimal winning bid of bidder 1 in A' , regardless of the realization of $v_j^1 \in [j \cdot \varepsilon, (j+1) \cdot \varepsilon)$, remains w_1 , and so for every $v_1 \in [0, H]$, we surely have that $r_1^{A'}(v_1, \dots, v_n) = r_1^A(v_1, \dots, v_n)$.

Proof of Part b(ii)(1): We reason by cases. For every $v_1 < j \cdot \varepsilon$, we have, regardless of the realization of $v_j^1 \in [j \cdot \varepsilon, (j+1) \cdot \varepsilon)$, that bidder 1 loses against v_{-1} in both A' and A . For every $v_1 \geq (j+1) \cdot \varepsilon$, we have, regardless of the realization of v_j^1 , that bidder 1 wins against v_{-1} in both A' and A .

Proof of Part b(ii)(2): The winning probability of a bid $v_1 \sim F_1|j$ in A is $\mathbb{P}(v_1 \geq w_1) = \mathbb{P}_{v_1 \sim F_1|j}(v_1 \geq w_1)$, and the winning probability of such a bid in A' is $\mathbb{P}(v_j^1 \geq w_1) = \mathbb{P}_{v_j^1 \sim F_1|j}(v_j^1 \geq w_1)$, i.e., the same probability, as required.

Proof of Part b(ii)(3): Since $w_1 \in (j \cdot \varepsilon, (j+1) \cdot \varepsilon)$, we have that the minimal winning bid w'_1 of bidder 1 in A' is surely either $j \cdot \varepsilon$ or $(j+1) \cdot \varepsilon$. Therefore, surely $|w'_1 - w_1| < \varepsilon$. \square

Proof of Lemma 3.5. By Lemma 3.3(a) the only revenue that is affected by the change from A to A' is the revenue from bidder i , and by Lemma 3.3(b)(i), this revenue is affected only when the minimal winning bid of bidder i lies in $(j \cdot \varepsilon, (j+1) \cdot \varepsilon)$, and even then, by Lemma 3.3(b)(ii), this revenue is affected with the probability that bidder i wins, and even then it is changed by less than an additive ε . So, the change in the overall revenue is less than an additive $\varepsilon \cdot p_j^i$, as required. \square

Proof of Lemma 3.6. Assume without loss of generality that $i = 1$. Let A'' be the result of applying to A (in arbitrary order) all F_1 -randomized ε -rounding actions on all of bidder 1's ε -intervals. We start by noting that by Lemma 3.3(a),

$$\mathbb{E}_{v_1 \sim F_1|_j} \left[\mathbb{E}_{A'} r_1^{A'}(v_1) \right] = \mathbb{E}_{v_1 \sim F_1|_j} \left[\mathbb{E}_{A''} r_1^{A''}(v_1) \right].$$

It therefore suffices to show that for every fixed $v_2, \dots, v_n \in [0, H]$, it is the case that

$$\left| \mathbb{E}_{v_1 \sim F_1|_j} \left[\mathbb{E}_{A''} r_1^{A''}(v_1, \dots, v_n) \right] - \mathbb{E}_{v_1 \sim F_1|_j} r_1^A(v_1, \dots, v_n) \right| < \varepsilon \cdot \mathbb{P}_{v_1 \sim F_1|_j} [1 \text{ wins in } A \text{ against } v_{-1}].$$

Let w_1 be the minimal winning bid of bidder 1 in A when the other bidders bid v_2, \dots, v_n . Let k be such that $w_1 \in [k \cdot \varepsilon, (k+1) \cdot \varepsilon)$, and let A''' be the result of applying the F_1 -randomized ε -rounding action on bidder 1's ε -interval k to A . By thinking of A'' as having been constructed from A by first applying the randomized rounding action for k to yield A''' (following which, the minimal winning bid of bidder 1 surely becomes either $k \cdot \varepsilon$ or $(k+1) \cdot \varepsilon$) and then applying all randomized rounding actions for intervals other than k , we obtain by Lemma 3.3(b)(i) that

$$\mathbb{E}_{v_1 \sim F_1|_j} \left[\mathbb{E}_{A''} r_1^{A''}(v_1, \dots, v_n) \right] = \mathbb{E}_{v_1 \sim F_1|_j} \left[\mathbb{E}_{A'''} r_1^{A'''}(v_1, \dots, v_n) \right].$$

We conclude the proof as by Lemma 3.3(b)(ii) we obtain that

$$\left| \mathbb{E}_{v_1 \sim F_1|_j} \left[\mathbb{E}_{A'''} r_1^{A'''}(v_1, \dots, v_n) \right] - \mathbb{E}_{v_1 \sim F_1|_j} r_1^A(v_1, \dots, v_n) \right| < \varepsilon \cdot \mathbb{P}_{v_1 \sim F_1|_j} [1 \text{ wins in } A \text{ against } v_{-1}].$$

(We note that when $w_1 = k \cdot \varepsilon$, Lemma 3.3(b)(i) implies that there is surely no difference between the two revenues above, and so Lemma 3.6 can in fact be slightly strengthened, by replacing p in the statement of this lemma with the probability that i wins in A and pays a price that is not an integer multiple of ε when bidding $v_i \sim F_i|_j$, when the remaining bids are distributed according to F_{-i} .) \square

Proof of Theorem 3.4. By the triangle inequality and by Lemma 3.6,

$$\left| \mathbb{E}_{A'} \text{Rev}^{A'}(F) - \text{Rev}^A(F) \right| \leq \sum_{i=1}^n \left| \mathbb{E}_{v_i \sim F_i} \left[\mathbb{E}_{A'} r_i^{A'}(v_i) - r_i^A(v_i) \right] \right| < \sum_{i=1}^n \varepsilon \cdot \mathbb{P}_F^i [i \text{ wins}] = p \cdot \varepsilon,$$

as required.

Alternatively, we may prove the result via Lemma 3.5: We start with A and gradually apply, for all i and j , the F_i -randomized ε -rounding action on bidder i 's interval j , to obtain A' . By Lemma 3.5, when the randomized rounding action for a pair (i, j) is applied, the revenue changes by less than an additive $p_j^i \cdot \varepsilon$, and so less than an additive $\varepsilon \cdot \mathbb{P}_F^i [i \text{ wins}]$ for all randomized rounding actions for each bidder i , for a total of less than an additive $p \cdot \varepsilon$ for the entire randomized rounding rule. (Note that p_j^i , in expectation over all randomized rounding actions already applied to bidders other than i , does not change during the application of the randomized rounding action for any bidder-index pair other than (i, j) .)

(We note that regardless of whether we prove Theorem 3.4 via Lemma 3.5 or via Lemma 3.6,²⁷ Theorem 3.4 can in fact be slightly strengthened, by replacing p in the statement of this theorem with the probability that some bidder wins in A and pays a price that is not an integer multiple of ε when the profile of bids is distributed according to F .) \square

²⁷For the latter, by the note concluding its proof.

B.2 Proofs for Section 5

Proof of Proposition 5.3. We will show that with probability at least $1 - \delta$, it is the case that $|\mathbb{E}_{\hat{F}} r - \mathbb{E}_F r| < \varepsilon$ for any fixed random variable $r : [0, H]^n \rightarrow [0, H]$. The proposition will follow by setting $r \triangleq r^A$.

Let $(v_1^1, \dots, v_n^1), \dots, (v_1^t, \dots, v_n^t)$ be t independently drawn random tuples sampled from F . Let $k_1 \triangleq 1$ and for every $i \in \{2, 3, \dots, n\}$, fix $k_i \in \{1, \dots, t\}$. We note that $(v_1^{k_1+j}, \dots, v_n^{k_n+j})$ for $j = 1, \dots, t$ (where addition of k_i and j wraps around from t to 1) are t independent random samples from $F_1 \times \dots \times F_n = F$. Therefore, letting $r_{k_2, \dots, k_n} \triangleq \frac{1}{t} \sum_{j=1}^t r(v_1^{k_1+j}, \dots, v_n^{k_n+j})$, we have by the Chernoff-Hoeffding Inequality that

$$\mathbb{P}\left(|r_{k_2, \dots, k_n} - \mathbb{E}_F r| \geq \varepsilon\right) \leq 2 \exp\left(-\frac{2t\varepsilon^2}{H^2}\right).$$

Noting that by definition $\mathbb{E}_{\hat{F}} r = \frac{1}{t^{n-1}} \sum_{k_2, \dots, k_n \in \{1, \dots, t\}} r_{k_2, \dots, k_n}$, taking the union bound, we have that

$$\mathbb{P}\left(|\mathbb{E}_{\hat{F}} r - \mathbb{E}_F r| \geq \varepsilon\right) \leq t^{n-1} \cdot 2 \exp\left(-\frac{2t\varepsilon^2}{H^2}\right).$$

We conclude the proof as there exists $t = \text{poly}(H, n, 1/\varepsilon, \log 1/\delta)$ such that

$$t^{n-1} \cdot 2 \exp\left(-\frac{2t\varepsilon^2}{H^2}\right) \leq \delta. \quad \square$$

Proof of Lemma 5.5. Let $\delta' \triangleq \frac{\delta}{|S_\varepsilon^n|+1}$. By Lemma 5.2, $|S_\varepsilon^n| \leq \exp(\text{poly}(H, n, 1/\varepsilon))$. Therefore, by Proposition 5.3 there exists $t = \text{poly}(H, n, 1/\varepsilon, \log 1/\delta') = \text{poly}(H, n, 1/\varepsilon, \log 1/\delta)$ such that for every $A \in S_\varepsilon^n \cup \{OPT(F)\}$ separately, with probability at least $1 - \delta'$, we have that

$$|\text{Rev}^A(F) - \text{Rev}^A(\hat{F})| < \varepsilon.$$

By the union bound, we therefore have that with probability at least $1 - (|S_\varepsilon^n| + 1) \cdot \delta' = 1 - \delta$, this holds for *all* auctions $A \in S_\varepsilon^n \cup \{OPT(F)\}$ simultaneously, as required. \square

B.3 Proofs for Section 6

Proof sketch of Theorem 6.3. Similar to the proof of Theorem 3.4. Roughly speaking, instead of having a revenue change of less than ε whenever bidder i wins and her minimal winning bid is inside the open ε -interval to which the rounding action is applied (as in Lemma 3.3), it is the case that if f is the fraction of the winnings of bidder i whose threshold prices are inside this ε -interval, then the revenue change is less than $f' \cdot \varepsilon$, where f' is the “part of f ” that i wins. \square

Proof of Proposition 6.4. Analogous to the proof of Proposition 3.7, with Theorem 3.4 replaced with Theorem 6.3. \square

Proof of Lemma 6.6. Let $A \triangleq ((\phi_i)_{i \in N}; \text{MAX}_X)$. The algorithm is as follows, for $D, E = \text{poly}(H, 1/\varepsilon, \log 1/\delta)$ that we define below:

1. For every $e \in \{1, \dots, E\}$:
 - Draw a valuation profile $(v_1^e, \dots, v_n^e) \sim \hat{F}$.

2. For every $d \in \{1, \dots, D\}$:
 - (a) Draw a (deterministic) ε -rounding rule from (the distribution specified by) the \hat{F} -randomized ε -rounding rule.
 - (b) Apply this ε -rounding rule to A to obtain an (ε -coarse) auction A'_d .
 - (c) Initialize $r_d \leftarrow 0$.
 - (d) For every $e \in \{1, \dots, E\}$:
 - Update²⁸ $r_d \leftarrow r_d + \frac{r^{A'_d}(v_1^e, \dots, v_n^e)}{E}$.
3. Output (the ironed virtual valuations of) the auction A'_d for $d = \arg \max_{d \in \{1, \dots, D\}} r_d$.

We now choose D and E . We would like to choose D such that with high probability, at least one of the drawn auctions A'_d has revenue not significantly lower than the expected revenue of the \hat{F} -randomized ε -rounding of A , which we denote by A' . More accurately, we would like D to satisfy the following:

$$\mathbb{P} \left(\frac{1}{D} \sum_{d=1}^D \text{Rev}^{A'_d}(\hat{F}) \leq \mathbb{E}_{A'} \text{Rev}^{A'}(\hat{F}) - \varepsilon/3 \right) \leq \frac{\delta}{2}. \quad (1)$$

By the Chernoff-Hoeffding Inequality, we have that

$$\mathbb{P} \left(\frac{1}{D} \sum_{d=1}^D \text{Rev}^{A'_d}(\hat{F}) \leq \mathbb{E}_{A'} \text{Rev}^{A'}(\hat{F}) - \varepsilon/3 \right) \leq \exp \left(-\frac{2D\varepsilon^2}{9H^2} \right).$$

Therefore, by choosing $D \triangleq \frac{9H^2}{2\varepsilon^2} \log \frac{2}{\delta} = \text{poly}(H, 1/\varepsilon, \log 1/\delta)$, we have that Eq. (1) is satisfied.

We would now like to choose E such that with high probability, the estimated revenue r_d for each A'_d is close to the true revenue of A'_d from \hat{F} . More accurately, we would like E to satisfy the following, for every $d \in \{1, \dots, D\}$:

$$\mathbb{P} \left(|r_d - \text{Rev}^{A'_d}(\hat{F})| \geq \varepsilon/3 \right) \leq \frac{\delta}{2D}. \quad (2)$$

By the Chernoff-Hoeffding Inequality, we have that

$$\mathbb{P} \left(|r_d - \text{Rev}^{A'_d}(\hat{F})| \geq \varepsilon/3 \right) \leq 2 \exp \left(-\frac{2E\varepsilon^2}{9H^2} \right).$$

Therefore, by choosing $E \triangleq \frac{9H^2}{2\varepsilon^2} \log \frac{4D}{\delta} = \text{poly}(H, 1/\varepsilon, \log 1/\delta)$, we have that Eq. (2) is satisfied for every $d \in \{1, \dots, D\}$.

Having chosen D and E , we now prove the correctness of the algorithm. By taking the union bound over Eq. (1) and over Eq. (2) for all $d \in \{1, \dots, D\}$, we obtain that with probability at least $1 - \delta$, both of the following hold:

- For every $d \in \{1, \dots, D\}$ it is the case that $|r_d - \text{Rev}^{A'_d}(\hat{F})| < \varepsilon/3$, and

²⁸The explicit formula of Myerson (1981) for the (truthful) payments in a Myersonian auction yields that given an ε -coarse Myersonian auction A for a tractable single-parameter environment, and given a valuation profile $(v_1, \dots, v_n) \in \text{supp } \hat{F}$, it is possible in our setting to deterministically compute not only the allocation of A in time $\text{poly}(H, n)$, but also the associated payments in time $\text{poly}(H, n, t)$.

- There exists $d' \in \{1, \dots, D\}$ such that $\text{Rev}^{A'_{d'}}(\hat{F}) > \mathbb{E}_{A'} \text{Rev}^{A'}(\hat{F}) - \varepsilon/3$.

Let $d \in \{1, \dots, D\}$ be such that r_d is maximal (so the algorithm outputs A'_d). In particular, $r_d \geq r_{d'}$. Therefore,

$$\begin{aligned} \text{Rev}^{A'_d}(\hat{F}) &> r_d - \varepsilon/3 \geq r_{d'} - \varepsilon/3 > \text{Rev}^{A'_{d'}}(\hat{F}) - 2 \cdot \varepsilon/3 > \\ &> \mathbb{E}_{A'} \text{Rev}^{A'}(\hat{F}) - 3 \cdot \varepsilon/3 = \mathbb{E}_{A'} \text{Rev}^{A'}(\hat{F}) - \varepsilon > \text{Rev}^A(\hat{F}) - (W_X + 1) \cdot \varepsilon, \end{aligned}$$

where the last inequality is by Theorem 6.3.

Finally, we estimate the number of random bits used by the algorithm. To draw E valuation profiles, the algorithm requires $O(E \cdot n \cdot \log t)$ random bits. To draw D roundings of A' , the algorithm requires $O(D \cdot n \cdot H/\varepsilon \cdot \log t)$ random bits. Therefore, in total the algorithm uses at most $\text{poly}(H, n, 1/\varepsilon, \log 1/\delta, \log t)$ random bits, as required. \square

Proof of Proposition 6.7. Let b be as in Lemma 6.6 for (the given) t and ε , and for guaranteeing success with probability at least $1 - \delta/2$. We start by finding a probability η such that both of the following hold:

- On the one hand, if for every $i \in N$ the distribution F_i has an atom $v_i \in [0, H]$ with probability $F_i(v_i) \geq 1 - \eta$, then the auction that extracts the entire social welfare of the outcome that maximizes the social welfare for (v_1, \dots, v_n) , maximizes the overall revenue from F up to an additive ε . Furthermore, in this case with probability at least $1 - \delta$ each such atom v_i can be recovered by taking the most common sample value out of $s_v = \text{poly}(\log n, \log 1/\delta)$ (chosen below) independent samples from F_i .
- On the other hand, if the above does not hold (i.e., if there exists $i \in N$ such that for every $v_i \in [0, H]$ it is the case that $F_i(v_i) < 1 - \eta$), then one can with probability at least $1 - \delta/2$ extract b random bits from $s_b = \text{poly}(H, n, 1/\varepsilon, \log 1/\delta, \log t)$ (chosen below) independent pairs of (independent) samples from F .

Assume first that for every $i \in N$ there exists $v_i \in [0, H]$ such that $F_i(v_i) \geq 1 - \eta$. Taking the union bound, we have that $F(v_1, \dots, v_n) \geq 1 - n\eta$. For every $i \in N$, let

$$\phi'_i(v'_i) \triangleq \begin{cases} v_i & v'_i \geq v_i, \\ -\infty & v'_i < v_i. \end{cases} \quad (3)$$

We note that

$$r^{((\phi'_i)_{i \in N}; \text{MAX}_X)}(v_1, \dots, v_n) = \sum_{i=1}^n x_i \cdot v_i \geq r^{\text{OPT}(F; X)}(v_1, \dots, v_n), \quad (4)$$

where $(x_1, \dots, x_n) \in X$ is the outcome that maximizes the social welfare $\sum_{i=1}^n x_i \cdot v_i$ for (v_1, \dots, v_n) . (The inequality in Eq. (4) holds as the revenue can never exceed the social welfare.) Therefore, by definition of (v_1, \dots, v_n) we have that

$$\text{Rev}^{((\phi'_i)_{i \in N}; \text{MAX}_X)}(F) \geq \text{Rev}^{\text{OPT}(F; X)}(F) - n\eta H.$$

Therefore, choosing any $\eta < \frac{\varepsilon}{nH}$, we obtain that if for every $i \in N$ there exists $v_i \in [0, H]$ such that $F_i(v_i) \geq 1 - \eta$, then $\text{Rev}^{((\phi'_i)_{i \in N}; \text{MAX}_X)}(F) > \text{Rev}^{\text{OPT}(F; X)}(F) - \varepsilon$, for $(\phi'_i)_{i \in N}$ as defined in Eq. (3). We choose $\eta \triangleq \frac{\varepsilon}{(n+3)H} < \frac{\varepsilon}{nH}$ as this also conveniently guarantees

that $\eta \leq 1/4$, which we will use later. (If $\varepsilon \geq H$, then any revenue is at most ε and the proposition is trivially satisfied.)

Still under the assumption that for every $i \in N$ there exists $v_i \in [0, H]$ such that $F_i(v_i) \geq 1 - \eta$, let freq_{v_i} be (a random variable that equals) the fraction of samples that equal v_i out of s_v independent samples from F_i . Recall that $\eta \leq 1/4$. By the Chernoff-Hoeffding Inequality, we have that

$$\mathbb{P}(\text{freq}_{v_i} \leq 1/2) \leq \mathbb{P}(\text{freq}_{v_i} \leq 1 - \eta - 1/4) \leq \exp\left(-\frac{s_v}{8}\right),$$

and so taking $s_v \triangleq \lceil 8 \log n / \delta \rceil = \text{poly}(\log n, \log 1/\delta)$ guarantees that $\mathbb{P}(\text{freq}_{v_i} > 1/2) \geq 1 - \delta/n$, and by the union bound we have that with probability at least $1 - \delta$, for every $i \in N$ the most common sample value out of s_v independent samples from F_i is v_i .

Assume now that there exists $i \in N$ such that for every $v_i \in [0, H]$ it is the case that $F_i(v_i) < 1 - \eta$. We now show that there exists $s_b = \text{poly}(H, n, 1/\varepsilon, \log 1/\delta, t, b)$ such that with probability at least $1 - \delta/2$, from s_b independent pairs of samples from F_i , at least b pairs are of two nonidentical samples. Since for every $v_i \in [0, H]$ it is the case that $F_i(v_i) < 1 - \eta$, there exists a partition of $[0, H]$ into two disjoint sets V, W such that $F_i(V) < 1 - \eta$ and $F_i(W) < 1 - \eta$. As we will choose $s_b \triangleq s' \cdot b$, by the union bound we upper bound the probability that from s independent pairs of samples from F_i less than b pairs are of two nonidentical samples, by b times the probability that s' independent pairs of samples from F_i are all of identical samples. We upper bound the latter probability by the probability that $2s'$ independent samples from F_i lie entirely in one of the sets V, W , which is less than $2(1 - \eta)^{2s'}$. Choosing the smallest integer s' such that $2(1 - \eta)^{2s'} \leq \frac{\delta}{2b}$ therefore guarantees that the probability that from s_b independent pairs of samples from F less than b pairs are of two nonidentical samples, is at most $\delta/2$. It therefore remains to estimate s' (and hence s_b). By definition, s' is the smallest integer such that $2s' \cdot \log(1 - \eta) \leq \log \frac{\delta}{4b}$, and so

$$\begin{aligned} s' &= \left\lceil \frac{\log \frac{\delta}{4b}}{2 \log(1 - \eta)} \right\rceil \leq \left\lceil \frac{\log \frac{\delta}{4b}}{-2\eta} \right\rceil = \left\lceil \frac{\log \frac{4b}{\delta}}{2\eta} \right\rceil = \left\lceil \frac{\log(4b) + \log 1/\delta}{2\eta} \right\rceil = \\ &= \left\lceil \frac{(n+3)H \cdot (\log(4b) + \log 1/\delta)}{2\varepsilon} \right\rceil = \text{poly}(H, n, 1/\varepsilon, \log 1/\delta, \log b), \end{aligned}$$

and so $s_b = s' \cdot b = \text{poly}(H, n, 1/\varepsilon, \log 1/\delta, \log t)$.

We are now ready to present the deterministic “ ε -rounding” algorithm. Let $s \triangleq \text{Max}\{2s_b, s_v, t\} = \text{poly}(H, n, 1/\varepsilon, \log 1/\delta, t)$. The algorithm is as follows:

1. Split the given s samples from F into $\lfloor s/2 \rfloor$ pairs of samples.
2. If less than b of the pairs are of two nonidentical samples:
 - (a) For each $i \in N$, let v_i be the most common out of the i th-coordinate values of the s samples.
 - (b) Output $e = 0$ and the functions $(\phi'_i)_{i \in N}$ where for every i , ϕ'_i is defined only for v_i , with $\phi'_i(v_i) \triangleq v_i$.²⁹

²⁹Recall that by Remark 2.8, as an ironed virtual valuation we interpret this ϕ'_i as coinciding with ϕ'_i as defined in Eq. (3).

3. Else:

- (a) From each of the first b of the pairs of samples that are of two nonidentical samples, compute a single (random) bit as follows: let i be minimal such that the i th coordinates of these two samples differ; the desired (random) bit is 1 if the i th coordinate of the first sample is smaller than the i th coordinate of the second sample, and 0 otherwise.
- (b) Deterministically run the algorithm from Lemma 6.6 on $(\phi_i)_{i \in N}$ using the b (random) bits computed above instead of the (at most) b random bits required by the algorithm.
- (c) Output $e = 1$ and the functions $(\phi'_i)_{i \in N}$ that were output by the algorithm from Lemma 6.6 when run in Step 3(b).

We start by noting that if the above algorithm outputs $e = 1$, then by Lemma 6.6, $((\phi'_i)_{i \in N}; \text{MAX}_X)$ is an ε -rounding of $((\phi_i)_{i \in N}; \text{MAX}_X)$. To prove that the algorithm is correct with probability at least $1 - \delta$, we first note that the algorithm can terminate in precisely one of two ways:

- 1. In Step 2(b), since less than b of the $\lfloor s/2 \rfloor$ pairs of samples are of two nonidentical samples.
- 2. In Step 3(c), since at least b of the $\lfloor s/2 \rfloor$ pairs of samples are of two nonidentical samples.

If the algorithm terminates in Step 3(c), then since the b bits computed in Step 3(a) are indeed random and independent (von Neumann, 1951), and so Lemma 6.6 guarantees that with probability at least $1 - \delta/2$, we have

$$\text{Rev}^{((\phi'_i)_{i \in N}; \text{MAX}_X)}(\hat{F}) > \text{Rev}^{((\phi_i)_{i \in N}; \text{MAX}_X)}(\hat{F}) - (W_X + 1) \cdot \varepsilon,$$

and therefore if the algorithm terminates in Step 3(c), then it succeeds with at least this probability. To complete the proof of correctness, we reason by cases, according to whether or not for every $i \in N$ there exists $v_i \in [0, H]$ such that $F_i(v_i) \geq 1 - \eta$.

We start by analyzing the case in which for every $i \in N$ there exists $v_i \in [0, H]$ such that $F_i(v_i) \geq 1 - \eta$. In this case, if the algorithm terminates in Step 2(b), then as shown above, with probability at least $1 - \delta$ we correctly identify v_i using $s \geq s_v$ samples, and so $e = 0$ and (as shown above) $\text{Rev}^{((\phi'_i)_{i \in N}; \text{MAX}_X)}(F) > \text{Rev}^{\text{OPT}(F; X)}(F) - \varepsilon$, as required. So, regardless of whether the algorithm terminates in Step 2(b) or in Step 3(c), with probability at least $1 - \delta$ the algorithm succeeds, as required.

It remains to analyze the case in which there exists $i \in N$ such that for every $v_i \in [0, H]$ it is the case that $F_i(v_i) < 1 - \eta$. In this case, as shown above, the probability that less than b of the $\lfloor s/2 \rfloor \geq s_b$ pairs of samples have two nonidentical i th coordinate-values is at most $\delta/2$. Therefore, with probability at least $1 - \delta/2$ the algorithm terminates in Step 3(c), and as shown above, if this happens, then with probability at least $1 - \delta/2$, the algorithm succeeds. Therefore, taking the union bound, in this case with probability at least $1 - \delta$ we have that $e = 1$ and $\text{Rev}^{((\phi'_i)_{i \in N}; \text{MAX}_X)}(\hat{F}) > \text{Rev}^{((\phi_i)_{i \in N}; \text{MAX}_X)}(\hat{F}) - (W_X + 1) \cdot \varepsilon$, as required. \square

Proof of Theorem 6.10. The proof is analogous to that of Theorem 5.6. An analogous version of Lemma 5.5 can be proven for S_ε^X rather than S_ε^n via an analogous proof, replacing Lemma 5.2 with the assumption that $|S_\varepsilon^X| \leq \exp(\text{poly}(H, n, 1/\varepsilon))$. (For position and

matroid environments, no replacement is needed, and for arbitrary deterministic environments the assumption holds by Lemma 6.9.) Let t be as in this analogue of Lemma 5.5 for $S_{\varepsilon/(W_X+3)}^X$, for guaranteeing success with probability at least $1 - \delta/2$. Let s be as in Proposition 6.7 for guaranteeing success with probability at least $1 - \delta/2$. By Theorem 2.7 and Proposition 6.7 (and by taking a union bound over two failure probabilities of at most $\delta/2$ each), all that we have to show is that if the algorithm of Proposition 6.7 outputs $e = 1$ and an $\frac{\varepsilon}{W_X+3}$ -rounded auction $A' \in S_{\varepsilon/(W_X+3)}^n$ that maximizes the revenue from \hat{F} up to less than an additive $(W_X + 1) \cdot \frac{\varepsilon}{W_X+3}$, then this auction also maximizes the revenue from F up to less than an additive ε . Showing this is analogous to the conclusion of the proof of Theorem 5.6. \square

Proof of Proposition 6.12. Analogous to the proof of Proposition 6.7 and Lemma 6.6, replacing every occurrence of MAX_X with APPROX_X , replacing every occurrence of $\text{Rev}^{\text{OPT}(F;X)}(F)$ with $\frac{\text{Rev}^{\text{OPT}(F;X)}(F)}{C}$, and with Eq. (4) becoming

$$r^{((\phi'_i)_{i \in N}; \text{APPROX}_X)}(v_1, \dots, v_n) = \sum_{i=1}^n x_i \cdot v_i \geq \frac{r^{\text{OPT}(F;X)}(v_1, \dots, v_n)}{C},$$

where $(x_1, \dots, x_n) \in X$ is the outcome that APPROX_X outputs given (v_1, \dots, v_n) , which maximizes the social welfare $\sum_{i=1}^n x_i \cdot v_i$ for (v_1, \dots, v_n) up to a multiplicative factor of C . \square

Proof of Theorem 6.14. Analogous (see Section 6.5 for details) to the proof of Theorem 6.10, with $S_{\varepsilon/(W_X+3)}^X$ replaced with $S_{\varepsilon/(W_X+3)}^{\text{APPROX}_X}$, and with Proposition 6.7 replaced with Proposition 6.12. \square

C Simple Efficient Empirical Revenue Maximization for I.I.D. Bidders

Possibly the most naïve way to construct a small set of auctions is to consider only auctions for which all of the parameters of the auction are specified with some fixed precision, and to learn an auction from this set for some empirical distribution by computing the optimal auction for that distribution, and then rounding its parameters to obtain an auction from this set. The best possible hope in this case would be that on each valuation profile for the bidders, the rounded auction loses less than an additive ε in revenue. While, as shown in Example A.1 in Appendix A, such an approach turns out too good to be true in the general case of non-i.i.d. product distributions, in this appendix we show that remarkably, this approach succeeds for arbitrary (even irregular) i.i.d. product distributions, reproving the existence of a polynomial-time learning algorithm from Roughgarden and Schrijvers (2016) via a considerably simpler argument and a more natural (albeit not dissimilar) algorithm, and highlighting the difference between the general case considered in the main text and the i.i.d. case.

C.1 Optimal Auctions

We once again make very weak use of Myerson's (1981) characterization of optimal auctions for i.i.d. product distributions, and once again only present this characterization,

which is a special case of that presented in Section 2, to the extent required by our analysis.

Definition C.1 (Second-Price Auction with Reserve Price and Ironed Intervals, Myerson, 1981). A *second-price auction with reserve price and ironed intervals* is a pair (p, I) , where $p \in [0, H]$ is called the *reserve price*, and $I \subseteq \{[\ell, h) \mid p \leq \ell < h \leq H\}$ are (possibly infinitely many) pairwise-disjoint intervals called *ironed intervals*. In this auction, there is a winner unless all bidders bid below the reserve price, and the winner is the bidder with lowest index among those with the highest bid, unless the highest bid lies in the same ironed interval as lower bids, in which case the winner is the bidder with lowest index among those whose bid lies in this ironed interval; the winner pays her minimal winning bid.³⁰

Myerson (1981) proved that for every continuous i.i.d. product distribution, there exists a second-price auction with reserve price and ironed intervals that obtains the optimal revenue. Elkind (2007) showed the same for discrete i.i.d. product distributions, giving an efficient algorithm for computing this optimal auction. These results are special cases of Theorems 2.6 and 2.7, respectively.

Theorem C.2 (Myerson, 1981). *For every distribution $F \in \Delta([0, H])$, there exists a second-price auction with reserve price and ironed intervals that for every $n \in N$ achieves maximal revenue from F^n among all possible auctions.*

Theorem C.3 (Elkind, 2007). *Let $t \in \mathbb{N}$. There exists an algorithm that runs in time $\text{poly}(t)$, such that given a discrete distribution $\hat{F} \in \Delta([0, H])$ with support of size at most t , outputs a second-price auction with reserve price and ironed intervals that for every $n \in N$ achieves maximal revenue from \hat{F}^n among all possible auctions.*

The reader may verify that a Myersonian auction $(\phi_i)_{i \in N}$ where $\phi_i = \phi$ for some function ϕ for all $i \in N$, coincides with a second-price auction with reserve price and ironed intervals, where the reserve price is $\phi^{-1}(0)$, and the ironed intervals are the intervals on which ϕ is constant and nonnegative.

C.2 Natural Efficient Distribution-Independent Rounding of Second-Price Auctions with Reserve Prices

As already noted above, we will now show that rounding the parameters that define a second-price auction with reserve price and ironed intervals to some fixed precision results in a loss in revenue of at most that precision on each valuation profile.

Definition C.4 (ε -Rounded-Down Auction). Let $\varepsilon > 0$.

- For every $x \in \mathbb{R}_+$, let $\lfloor x \rfloor_\varepsilon \triangleq \varepsilon \cdot \lfloor x/\varepsilon \rfloor$ be the value of x , rounded down to the nearest integer multiple of ε .

³⁰Once again, this auction can also be made symmetric, by choosing the winner uniformly at random among all bidders with highest bid / whose bids lie in the same ironed interval as the highest bid, and adapting the payments accordingly. The revenue is the same either way. As noted above, in this paper we use lexicographic ordering for simplicity.

- For every second-price auction with reserve price and ironing intervals (p, I) , we define its ε -rounded-down counterpart $\lfloor(p, I)\rfloor_\varepsilon$ as follows:

$$\lfloor(p, I)\rfloor_\varepsilon \triangleq \left(\lfloor p \rfloor_\varepsilon, \left\{ \lfloor \ell \rfloor_\varepsilon, \lfloor h \rfloor_\varepsilon \mid [\ell, h] \in I \text{ \& } \lfloor \ell \rfloor_\varepsilon < \lfloor h \rfloor_\varepsilon \right\} \right).$$

Lemma C.5. *For every (p, I) , for every $\varepsilon > 0$, for every $n \in \mathbb{N}$, and for every $v_1, \dots, v_n \in [0, H]$, it is the case that*

$$r^{\lfloor(p, I)\rfloor_\varepsilon}(v_1, \dots, v_n) > r^{(p, I)}(v_1, \dots, v_n) - \varepsilon.$$

The proof of Lemma C.5 and of other results from this appendix are relegated to Appendix C.4. Lemma C.5 implies that the loss in revenue is less than an additive ε also on every distribution.

Proposition C.6. *For every (p, I) , for every $\varepsilon > 0$, for every $F \in \Delta([0, H])$ and for every $n \in \mathbb{N}$, it is the case that*

$$\text{Rev}^{\lfloor(p, I)\rfloor_\varepsilon}(F^n) > \text{Rev}^{(p, I)}(F^n) - p \cdot \varepsilon,$$

where p is the probability that some bidder wins in (p, I) when the profile of bids is distributed according to F^n .³¹

We note that Lemma C.5 and Proposition C.6 (and therefore also all of the results of Appendix C.3) would still hold (via similar analysis) even if we were to ε -round $\lfloor(p, I)\rfloor_\varepsilon$ into an ε -coarse auction by rounding its inputs (in addition to the parameters of the auction) down to the nearest integer multiple of ε before applying its allocation rule (while adapting the payments accordingly). We further note that a construction similar to that of Example A.3 in Appendix A shows that the bound of “less than ε ” on the revenue loss in Proposition C.6 (and hence also in Lemma C.5) cannot be unconditionally tightened any further, even if $\lfloor(p, I)\rfloor_\varepsilon$ is replaced with an arbitrary ε -coarse auction, and already for $n = 1$ bidder (more bidders with the same valuation distribution may be added to show the same result for $n > 1$ bidders).

C.3 Uniform Convergence over the Set of Rounded Second-Price Auctions with Reserve Prices

Analogously to the main text, we obtain a uniform convergence result through showing that the set of ε -rounded-down second-price auctions with reserve price and ironing intervals is small, and then utilize this and the previous results of this appendix to reprove the analogous learning result for i.i.d. produce distributions.

Definition C.7 ($\lfloor S \rfloor_\varepsilon$). For every $\varepsilon > 0$, we denote the set of all ε -rounded-down second-price auctions with reserve price and ironing intervals with parameters in $[0, H]$ by $\lfloor S \rfloor_\varepsilon$.

Lemma C.8. $|\lfloor S \rfloor_\varepsilon| \leq \exp(\text{poly}(H, 1/\varepsilon)).$

³¹Interestingly, similarly (at least conceptually) to Proposition 3.7 (and Proposition 6.4), a slightly stronger statement than that of Proposition C.6 also holds, where p is replaced with the probability that some bidder wins in A and pays a price that is not an integer multiple of ε when the profile of bids is distributed according to F . The case analysis is similar to the one in the proof of Lemma C.5, and is left to the reader.

Proof. To specify the ironing intervals, it is more than sufficient to merely specify, for each multiple of ε in $[p, H]$, whether an ironed interval ends at that point, starts at that point, neither, or both. \square

To prove an analogous result to Theorem 5.6/Theorem 1.4 for i.i.d. distributions using the results of Appendix C.2, we require the following analogue of Lemma 5.5.

Lemma C.9. *For every $\varepsilon > 0$ and $\delta > 0$, there exists $t = \text{poly}(H, n, 1/\varepsilon, \log 1/\delta)$ such that the following holds. Fix $F \in \Delta([0, H])$, draw t samples from F , and let \hat{F} be the empirical uniform distribution over these t samples. With probability at least $1 - \delta$, it is the case that*

$$|\text{Rev}^A(F^n) - \text{Rev}^A(\hat{F}^n)| < \varepsilon$$

holds simultaneously for every $A \in \lfloor S \rfloor_\varepsilon \cup \{OPT(F^n)\}$.

The proof of Lemma C.9 is analogous to that of Lemma 5.5 (taking a union bound due to the exponential size of $\lfloor S \rfloor_\varepsilon$), with Proposition 5.3 replaced with the following analogous proposition for i.i.d. distributions. We phrase and prove the following proposition as a general probabilistic concentration inequality, which we also find interesting in its own right. (When reading this proposition, the reader is encouraged to think of the random variable r as r^A for some fixed auction A .)

Proposition C.10. *For every $\varepsilon > 0$ and $\delta > 0$, there exists $t = \text{poly}(H, n, 1/\varepsilon, \log 1/\delta)$ such that the following holds. Fix $F \in \Delta([0, H])$ and fix $r : [0, H]^n \rightarrow [0, H]$. Draw t samples from F , and let \hat{F} be the empirical uniform distribution over these t samples. Then, with probability at least $1 - \delta$ it is the case that*

$$\left| \mathbb{E}_{\hat{F}^n} r - \mathbb{E}_{F^n} r \right| < \varepsilon.$$

Remark C.11. We note that Proposition C.10 does not follow from Proposition 5.3 (or from the more general concentration inequality of Devanur et al., 2016). Indeed, while in the latter a random tuple from $\hat{F}_1 \times \cdots \times \hat{F}_n$ distributes according to $F_1 \times \cdots \times F_n$, in the former a random tuple from \hat{F}^n does not necessarily distribute according to F^n . Indeed, a random tuple from \hat{F}^n has, e.g., positive probability for containing duplicate values, while a random tuple from F^n may have (depending on F) zero probability for containing duplicate values.

Remark C.12. Concentration inequalities similar to Propositions 5.3 and C.10 can be similarly shown to hold for cases in which the n distributions (e.g., over valuations of bidders) are divided into subsets, where in each subset the distributions are i.i.d.

Combining Lemma C.9, Proposition C.6, and Theorem C.3, we obtain the following analogue of Theorem 5.6/Theorem 1.4, providing a natural polynomial-time algorithm for learning an approximately optimal auction from samples from an arbitrary unknown bounded i.i.d. product distribution.

Theorem C.13. *There exists $t = \text{poly}(H, n, 1/\varepsilon, \log 1/\delta)$ such that the following holds. Let F be an arbitrary distribution on $[0, H]$. Draw t samples from F , and let \hat{F} be the empirical distribution over these t samples. Then, with probability at least $1 - \delta$, the optimal auction for \hat{F}^n (which can be deterministically computed in time $\text{poly}(t)$ via Theorem C.3), where the reserve price as well as the bounds of each ironed interval are all rounded down to the nearest multiple of $\varepsilon/3$, approximates the maximum possible revenue from F^n up to less than an additive ε .*

C.4 Proofs for Appendix C

Proof Lemma C.5. For a value $v \in [0, H]$, define

$$h_{(p,I)}(v) \triangleq \begin{cases} p & v \leq p \\ h & \exists [\ell, h) \in I : v \in [\ell, h) \\ v & \text{otherwise,} \end{cases}$$

and

$$\ell_{(p,I)}(v) \triangleq \begin{cases} p & v \leq p \\ \ell & \exists [\ell, h) \in I : v \in [\ell, h) \\ v & \text{otherwise,} \end{cases}$$

and note that for a valuation profile with second-highest bid $v^{(2)}$, the revenue from this profile, assuming that not all bids are lower than the reserve price p , is either $h_{(p,I)}(v^{(2)})$ or $\ell_{(p,I)}(v^{(2)})$.

Let $(v_1, \dots, v_n) \in [0, H]^n$ be a valuation profile. To show that $r^{(p,I)}(v_1, \dots, v_n) < r^{\lfloor(p,I)\rfloor_\varepsilon}(v_1, \dots, v_n) + \varepsilon$, we reason by cases. We note that if the revenue from (p, I) is zero, then there is nothing to prove. We assume therefore that the highest bid is no less than p , and so also no less than $\lfloor p \rfloor_\varepsilon$.

- If the winning bid in (p, I) lies in the same ironed interval (in (p, I)) as another bid:
Let $v^{(2)}$ be the second-highest bid (which may or may not be the winning bid). In this case, $r^{\lfloor(p,I)\rfloor_\varepsilon}(v_1, \dots, v_n) \geq \ell_{\lfloor(p,I)\rfloor_\varepsilon}(v^{(2)}) > \ell_{(p,I)}(v^{(2)}) - \varepsilon = r^{(p,I)}(v_1, \dots, v_n) - \varepsilon$.
- If the winning bid in (p, I) does not lie in the same ironed interval (in (p, I)) as another bid:

Let $v^{(2)}$ be the second-highest bid.

- If the winning bid in $\lfloor(p, I)\rfloor_\varepsilon$ lies in the same ironed interval (in $\lfloor(p, I)\rfloor_\varepsilon$) as another bid:

This means that $v^{(2)}$ lies inside an ironed interval that “before rounding” did not contain it. Therefore (regardless of whether or not the winning bidder changed), $r^{\lfloor(p,I)\rfloor_\varepsilon}(v_1, \dots, v_n) \geq \ell_{\lfloor(p,I)\rfloor_\varepsilon}(v^{(2)}) = \lfloor v^{(2)} \rfloor_\varepsilon > h_{(p,I)}(v^{(2)}) - \varepsilon \geq r^{(p,I)}(v_1, \dots, v_n) - \varepsilon$, as required.

- If the winning bid in $\lfloor(p, I)\rfloor_\varepsilon$ does not lie in the same ironed interval (in $\lfloor(p, I)\rfloor_\varepsilon$) as another bid:

Recall that the revenue from (p, I) is $\ell_{(p,I)}(v^{(2)})$ if $v^{(2)}$ lies in an ironed interval (in (p, I)) containing bids only from bidders with indices higher than that of the winner, and otherwise $h_{(p,I)}(v^{(2)})$. Similarly, the revenue from $\lfloor(p, I)\rfloor_\varepsilon$ is $\ell_{\lfloor(p,I)\rfloor_\varepsilon}(v^{(2)})$ if $v^{(2)}$ lies in an ironed interval (in $\lfloor(p, I)\rfloor_\varepsilon$) containing bids only from bidders with indices higher than that of the winner, and otherwise $h_{\lfloor(p,I)\rfloor_\varepsilon}(v^{(2)})$.

Since $h_{\lfloor(p,I)\rfloor_\varepsilon}(v^{(2)}) > h_{(p,I)}(v^{(2)}) - \varepsilon$ and $\ell_{\lfloor(p,I)\rfloor_\varepsilon}(v^{(2)}) > \ell_{(p,I)}(v^{(2)}) - \varepsilon$ (and since $h_{(p,I)}(v^{(2)}) \geq \ell_{(p,I)}(v^{(2)})$ and $h_{\lfloor(p,I)\rfloor_\varepsilon}(v^{(2)}) \geq \ell_{\lfloor(p,I)\rfloor_\varepsilon}(v^{(2)})$), we need only verify that no significant revenue loss occurred if the winner pays $h_{(p,I)}(v^{(2)}) > \ell_{(p,I)}(v^{(2)})$ in (p, I) but pays $\ell_{\lfloor(p,I)\rfloor_\varepsilon}(v^{(2)}) < h_{\lfloor(p,I)\rfloor_\varepsilon}(v^{(2)})$ in $\lfloor(p, I)\rfloor_\varepsilon$. In this case, some bid $v' < v^{(2)}$ (of a bidder with low index) that lies in the same

ironed interval in (p, I) as $v^{(2)}$ does not lie in the same ironed interval as $v^{(2)}$ in $\lfloor (p, I) \rfloor_\varepsilon$. Furthermore, in this case $v^{(2)}$ lies inside some ironed interval in $\lfloor (p, I) \rfloor_\varepsilon$. This means that $v^{(2)}$ lies inside an ironed interval in $\lfloor (p, I) \rfloor_\varepsilon$ that “before rounding” did not contain it. Therefore, $r^{\lfloor (p, I) \rfloor_\varepsilon}(v_1, \dots, v_n) = \ell_{\lfloor (p, I) \rfloor_\varepsilon}(v^{(2)}) = \lfloor v^{(2)} \rfloor_\varepsilon > h_{(p, I)}(v^{(2)}) - \varepsilon = r^{(p, I)}(v_1, \dots, v_n) - \varepsilon$, as required. \square

Proof of Proposition C.6. Immediate from Lemma C.5, noting that no revenue loss occurs for valuation profiles for which the auction (p, I) chooses no winner. \square

Proof of Proposition C.10. The main idea behind the proof is similar to the one behind that of Proposition 5.3, but keeping the samples independent is somewhat more delicate and involved. Let v^1, \dots, v^t be t independently drawn random values sampled from F . Let

$$K \triangleq \{(k_1, \dots, k_n) \mid k_1 = 1 \ \& \ \forall i \in \{2, \dots, n\} : k_i \in \{1, \dots, t\} \setminus \{k_1, \dots, k_{i-1}\}\},$$

and note that $|K| = \frac{(t-1)!}{(t-n)!}$. Fix $(k_1, \dots, k_n) \in K$. Since k_1, \dots, k_n are distinct, we note that for every $j \in \{1, \dots, t\}$, the tuple $(v^{k_1+j}, \dots, v^{k_n+j})$ (where addition of k_i and j wraps around from t to 1) is a random sample from F^n . We now greedily partition the set of t n -tuples of indices $I = I(k_1, \dots, k_n) \triangleq \{k_1 + j, k_2 + j, \dots, k_n + j\}_{j=1}^t$ into sets such that in each such set, the indices in all coordinates of all n -tuples are distinct. Let I_0 be a maximum subset of I where all indices are distinct, let I_1 be a maximum subset of $I \setminus I_0$ where all indices are distinct, let I_2 be a maximum subset of $I \setminus (I_0 \cup I_1)$ where all indices are distinct, etc.

We first claim that $I_0 \cup I_1 \cup \dots \cup I_{n \cdot (n-1)} = I$, i.e., that this process concludes after at most $n \cdot (n-1) + 1$ greedy steps. To show this, it is enough to show that $|I_i| \geq \frac{|I \setminus \{I_0, \dots, I_{i-1}\}|}{n \cdot (n-1) + 1 - i}$ for every $i = 0, \dots, n \cdot (n-1)$, as this implies that $I_{n \cdot (n-1)} = I \setminus \{I_0, \dots, I_{n \cdot (n-1)-1}\}$. Indeed, for each such i , each n -tuple in I_i “intersects” at most $n \cdot (n-1)$ other n -tuples in I (for every j , the j th coordinate of this n -tuple “intersects” the k th coordinate of precisely one other n -tuple in I , for $k \neq j$), at least one of which is contained in each of I_0, \dots, I_{i-1} (otherwise, this n -tuple could have been added to one of these sets), and so each n -tuple in I_i “blocks” at most $n \cdot (n-1) - i$ n -tuples from $I \setminus \{I_1, \dots, I_{i-1}\}$ from being added to I_i , obtaining that $|I_i| \geq \frac{|I \setminus \{I_0, \dots, I_{i-1}\}|}{n \cdot (n-1) + 1 - i}$ (under a worst-case scenario where the sets of “blocked” n -tuples are disjoint for every two “blocking” n -tuples), as required.

By construction, for every i , we have that $(v^{k_1+j}, \dots, v^{k_n+j})$ for all j such that $(k_1 + j, \dots, k_n + j) \in I_i$, are $|I_i|$ independent random samples from F^n . Therefore, for every i with $|I_i| \geq \sqrt{t}$, letting $r_{k_1, \dots, k_n}^i \triangleq \frac{1}{|I_i|} \sum_{(k_1+j, \dots, k_n+j) \in I_i} r(v^{k_1+j}, \dots, v^{k_n+j})$, we have by the Chernoff-Hoeffding Inequality that

$$\mathbb{P}\left(\left|r_{k_1, \dots, k_n}^i - \mathbb{E}_{F^n} r\right| \geq \varepsilon/2\right) \leq 2 \exp\left(-\frac{|I_i| \varepsilon^2}{2H^2}\right) \leq 2 \exp\left(-\frac{\sqrt{t} \varepsilon^2}{2H^2}\right).$$

Let $U(k_1, \dots, k_n) \triangleq \bigcup_{i: |I_i| \geq \sqrt{t}} I_i$. By definition, for $t \geq (n \cdot (n-1) + 1) \cdot \sqrt{t}$, there exists i such that $|I_i| \geq \sqrt{t}$ and so we have that $|U(k_1, \dots, k_n)| = |I| - |\bigcup_{i: |I_i| < \sqrt{t}} I_i| > t - n \cdot (n-1) \cdot \sqrt{t}$.

Let now $U \triangleq \bigcup_{(k_1, \dots, k_n) \in K} U(k_1, \dots, k_n)$. Taking the union bound, we have that

$$\mathbb{P}\left(\left|\left(\frac{1}{|U|} \cdot \sum_{\substack{(k_1, \dots, k_n) \in K \\ i: |I_i(k_1, \dots, k_n)| \geq \sqrt{t}}} |I_i| \cdot r_{k_1, \dots, k_n}^i\right) - \mathbb{E}_{F^n} r\right| \geq \varepsilon/2\right) \leq \frac{(t-1)!}{(t-n)!} \cdot (n \cdot (n-1) + 1) \cdot 2 \exp\left(-\frac{\sqrt{t} \varepsilon^2}{2H^2}\right).$$

Noting that by definition,

$$\mathbb{E}_{\hat{F}^n} r = \frac{|U|}{t^n} \cdot \overbrace{\frac{1}{|U|} \cdot \sum_{\substack{(k_1, \dots, k_n) \in K \\ i: |I_i(k_1, \dots, k_n)| \geq \sqrt{t}}} |I_i| \cdot r_{k_1, \dots, k_n}^i}^{\approx \mathbb{E}_{F^n} r} + \left(1 - \frac{|U|}{t^n}\right) \cdot \overbrace{\frac{1}{t^n - |U|} \cdot \sum_{\substack{(k_1, \dots, k_n) \in \\ \{1, \dots, t\}^n \setminus U}} r(v^{k_1}, \dots, v_n^k)}^{\leq H},$$

and noting that $|U| > |K| \cdot (t - n \cdot (n - 1) \cdot \sqrt{t}) = \frac{(t-1)!}{(t-n)!} \cdot (t - n \cdot (n - 1) \cdot \sqrt{t})$, we conclude the proof as there exists $t = \text{poly}(H, n, 1/\varepsilon, \log 1/\delta)$ such that both

$$\frac{(t-1)!}{(t-n)!} \cdot (n \cdot (n-1) + 1) \cdot 2 \exp\left(-\frac{\sqrt{t}\varepsilon^2}{2H^2}\right) \leq \delta \quad \text{and}$$

$$\frac{\frac{(t-1)!}{(t-n)!} \cdot (t - n \cdot (n-1) \cdot \sqrt{t})}{t^n} \geq 1 - \frac{\varepsilon}{2H}. \quad \square$$

Proof of Lemma C.9. Analogous to the proof of Lemma 5.5, with S_ε^n replaced with $\lfloor S \rfloor_\varepsilon$, and with Lemma 5.2 and Proposition 5.3 replaced with Lemma C.8 and Proposition C.10, respectively. \square

Proof of Theorem C.13. Analogous to the proof of Theorem 5.6, with $S_{\varepsilon/(n+2)}^n$ replaced with $\lfloor S \rfloor_{\varepsilon/3}$, and with Theorem 2.7, Proposition 4.3, and Lemma 5.5 replaced with Theorem C.3, Proposition C.6, and Lemma C.9, respectively. \square